EXHIBIT B

IN THE UNITED STATES DISTRICT COURT
WESTERN DISTRICT OF TEXAS
AUSTIN DIVISION

Intellectual Ventures I LLC and Intellectual Ventures II LLC,	
Plaintiffs/Counter-Defendants,	**Civil Action No. 1:19-CV-01075-ADA**
v.	**JURY TRIAL DEMANDED**
VMware, Inc.,	
Defendant/Counter-Plaintiff.	

DEFENDANT VMWARE, INC.'S OPENING CLAIM CONSTRUCTION BRIEF

TABLE OF CONTENTS

TABLE OF AUTHORITIES

Page(s)

Cases

Statutes

TABLE OF ABBREVIATIONS

Abbreviation	Full Name
PTO	United States Patent and Trademark Office
POSITA	Person of Ordinary Skill in the Art
Snoeren Decl.	Declaration of Alex Snoeren, Ph.D. Regarding Claim Construction
'686 patent	U.S. Patent No. RE 44,686
'726 patent	U.S. Patent No. RE 42,726
'937 patent	U.S. Patent No. 6,985,937
'937 FH	File History of U.S. Patent No. 6,985,937
'752 patent	U.S. Patent No. 7,949,752
'051 patent	U.S. Patent No. RE 43,051
'818 patent	U.S. Patent No. RE 44,818

TABLE OF EXHIBITS

Exhibit	Title
Ex. 1	Parties' Agreed Constructions to Claim Terms
Ex. 2	Declaration of Alex Snoeren, Ph.D. Regarding Claim Construction
Ex. 3	'937 Patent File History, Applicant Arguments dated November 17, 2003
Ex. 4	U.S. Patent No. 5,933,603 to Vahalia et al.
Ex. 5	'937 Patent File History, Non-Final Office Action dated May 10, 2004
Ex. 6	USPTO Patent Search of claim term "dynamic virtual server mover"
Ex. 7	Terms in the '726 patent that overlap with disputed claim terms in the '686 patent
Ex. 8	The parties' proposals for the terms in the '051 involving multiple recitations of "physical interface(s)" and "tunnel identifier(s)"
Ex. 9	'752 Patent File History, Response to Office Action dated March 16, 2009
Ex. 10	'752 Patent File History, Office Action dated May 5, 2009
Ex. 11	'752 Patent File History, Office Action dated October 30, 2009
Ex. 12	'752 Patent File History, Response to Office Action dated November 13, 2009
Ex. 13	'752 Patent File History, Office Action dated March 16, 2010
Ex. 14	'752 Patent File History, Response to Office Action dated August 20, 2010
Ex. 15	'752 Patent File History, Response to Office Action dated August 5, 2009
Ex. 16	'752 Patent File History, Final Office Action dated November 8, 2010
Ex. 17	'752 Patent File History, Response to Final Office Action dated January 4, 2011
Ex. 18	Report and Recommendation of United States Magistrate Judge, Intellectual Ventures v. HCC Ins. Holdings, Inc., Case No. 6:15-cv-660 (E.D. Tex. August 26, 2016) ("'752 Patent Magistrate Report")
Ex. 19	Charles Aulds, Linux Apache Web Server Administration, 39 (2001) ("Aulds")
Ex. 20	Barry Nusbaum, WebSphere Application Servers: Standard and Advanced Features 45 (1999)
Ex. 21	Ludmila Cherkasova, FLEX: Design and Management Strategy for Scalable Web Hosting Service, 14–15 (Oct. 1999)
Ex. 22	October 14, 2003 Amendment & Remarks, U.S. Patent Appl. No. 09/526,980
Ex. 23	Physical Interface, The IEEE Standard Dictionary of Electrical and Electronics Terms (6th ed., 1996)
Ex. 24	RE43,051 (U.S. Patent Appl. No. 11/858,091) Patent File History, May 17, 2010 Office Action
Ex. 25	U.S. Patent No. 6,286,047 ("Ramanathan")
Ex. 26	U.S. Patent No. 6,247,057 ("Barrera")
Ex. 27	Webster definition of consume
Ex. 28	'752 Patent List of Disputed Claim Terms
Ex. 29	BEN LAURIE AND PETER LAURIE, APACHE: THE DEFINITIVE GUIDE 163, 177, 242–43, 295 (1999)
Ex. 30	*Virtual Server*, MICROSOFT COMPUTER DICTIONARY (5th ed. 2002)
Ex. 31	'789 Patent (Orig. Patent) prosecution history, 2009.08.25 Resp. to Office Action at 12
Ex. 32	Competing Parties' Proposals for the '818 Patent
Ex. 33	Grotto Networking, available at https://www.grotto-networking.com/BBQoS.html
Ex. 34	http://tldp.org/HOWTO/Traffic-Control-HOWTO/index.html

Exhibit	Title
Ex. 35	https://lartc.org/howto/lartc.qdisc.classful.html#AEN1071
Ex. 36	Linux Advanced Routing & Traffic Control HOWTO, located at http://www.oamk.fi/~jukkao/lartc.pdf
Ex. 37	Traffic Control HOWTO, Version 1.0.2, Martin A. Brown, located at https://www.tldp.org/HOWTO/html_single/Traffic-Control-HOWTO/
Ex. 38	Traffic Control using tcng and HTB HOWTO, Version 1.0.1, Martin A. Brown, April 2006, located at http://linux-ip.net/articles/Traffic-Control-tcng-HTB-HOWTO.html
Ex. 39	Bavier, et al., Operating System Support for Planetary-Scale Network Services, Proceedings of the First Symposium on Networked Systems Design and Implementation (March 2004), located at https://www.usenix.org/legacy/events/nsdi04/tech/full_papers/bavier/bavier.pdf
Ex. 40	Benita, Kernel Korner - Analysis of the HTB Queuing Discipline Software, available at https://www.linuxjournal.com/article/7562 as of January 26, 2005 and printed in the Linux Journal, March 2005.
Ex. 41	U.S. Patent 7,161,904 titled: "System and method for hierarchical metering in a virtual router based network switch" to Hussein et al.
Ex. 42	Bavier et al, Container-based Operating SystemVirtualization: AScalable,High-performance Alternative to Hypervisors, Conference Paper in ACM SIGOPS Operating Systems Review, January 2007, located at http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.1018.1012&rep=rep1&type=pdf
Ex. 43	Valenzuela, J.L., et al., "A Hierarchical Token Bucket Algorithm to Enhance QoS in IEEE 802.11: Proposal, Implementation and Evaluation, IEEE, vol. 4, Sep. 2004 ("Valenzuela Article")
Ex. 44	Email from J. Deblois to M. Rueckheim dated March 3, 2020
Ex. 45	Webster's New Work Telecom Dictionary, Definition of Layer

I. INTRODUCTION

This case is immensely complicated. IV is asserting more than 40 claims, involving diverse technology from four separate patent families, that were developed by three companies acting independently. Indeed, the total number of words in the asserted claims exceeds 3,700, and hundreds of VMware products and features are accused of infringement. This is an unmanageable number of claims and issues for a jury to comprehend in a trial. Despite the significant number of issues, the parties have worked together to agree to constructions for certain terms (see Ex. 1) and have grouped the remaining disputes into the approximately 30 categories discussed herein. These remaining categories can primarily be summarized into five main groups.

First, for many technical terms, such as "hierarchical token bucket resource allocation" and "physical interface," VMware has proposed constructions consistent with the plain and ordinary meaning of the terms and consistent with the intrinsic record disclosure, while IV has simply proposed they be construed as "plain and ordinary meaning" or argued that the terms are "not amenable to construction." For these terms, IV has confirmed during the meet and confer process that it does not intend to offer alternative proposed constructions or interpretations. Thus, under *O2 Micro Int'l Ltd. v. Beyond Innovation Tech. Co*., 521 F.3d 1351 (Fed. Cir. 2008) the Court should adopt VMware's proposed constructions.

Second, for certain terms, like "modify a resource allocation" and "exhausted," the respective patent applicants significantly narrowed the claim scope during prosecution in response to PTO rejections. For example, the for term "exhausted"—the applicant actually argued that the prior art's art disclosure of the disclaimed scope (reusable resources) was "***the exact opposite***" of the amended claim language. VMware's proposals hold the applicants to their clear disclaimers consistent with the law. IV asks the Court to ignore these clear disclaimers in favor of an improperly broad reading of the claims.

Third, many terms involve Section 112 ¶ 6 means-plus-function considerations. For some of these terms, reciting classic means-for language, the parties agree that Section 112 ¶ 6 applies, but have disputes as to the specific structures. For some terms, there is no supporting structure counseling in favor of indefiniteness. For the others, VMware's proposed structures are properly limited to the specific portions of the specification that disclose the structures—while IV improperly proposes that the Court construe the structures with respect to ambiguous black-box type terminology. For others of these terms, the parties disagree whether Section 112 ¶ 6 applies. VMware's proposals for these terms recognize that the terms' usage of nonce terminology, like a "component," "module," or "program code" for performing a function, requires application of Section 112 ¶ 6 under *Williamson v. Citrix Online, LLC*, 792 F.3d 1339 (Fed. Cir. 2015). IV's proposal, in contrast, would leave these terms broad, despite the limited disclosure of corresponding structures in the specifications.

Fourth, for several disputes, for example the "virtual server" term, IV seeks to improperly change the meaning of technical terms to craft a better infringement read against VMware's products. Finally, for the remaining disputes, such as the various interfaces/tunnels terms in the '051 patent, and the "enforc[e/ing]", "receiv[e/ing]", etc. terms in the '818 patent, VMware's proposals address significant antecedent basis problems, or seek to clarify the complex technical terminology which may be confusing for a jury. IV's proposals reject these essential clarifications without explanation. VMware's proposals should be adopted at least for the reasons set forth herein and in the declaration of its expert Dr. Alex Snoeren submitted herewith.

II. DISPUTED TERMS FROM U.S. PATENT NO. RE44,686 (the "'686 patent")

A. "modif[y/ied] [a] resource allocation" / "modify[ing] [the] computer resources allocated to a virtual server" ('686 patent claims 5-7)

VMware Proposal	IV Proposal

"modif[y/ied] [a] quality of service guarantee" / "modify[ing] [the] quality of service guarantee of a virtual server" *See also* construction of "quality of service guarantee"	"modif[y/ied] set of functions and features of a physical host(s) used in implementing tasks for each virtual server" / "modify[ing] a set of the functions and features of a physical host(s) used in implementing tasks for each virtual server"

VMware's proposal is based on a clear and unmistakable disclaimer by the applicant during the prosecution of U.S. Patent No. 6,985,937 (the "'937 patent"), from which the '686 reissued. VMware's approach to these terms is also consistent with the claims, the specification,[1] and the applicant's own statements regarding the purpose of the claimed invention. By contrast, IV's proposal disregards the intrinsic record. Furthermore, rather than construe the phrase "modifying a *resource allocation*," IV's proposal reads "allocation" entirely out of the claim term and instead seeks to construe the phrase "modifying a resource."

During prosecution of the '937 patent, the examiner issued a rejection based on U.S. Patent No. 5,933,603 ("Vahalia"). Vahalia is generally directed to providing video-on-demand services by dynamically assigning resources to scheduled tasks. Ex. 4 (Vahalia, Abstract). In attempting to overcome the rejection, the applicant argued that "***[a] resource allocation for a process is specified as a quality of service guarantee. Thus, increasing a quality of service guarantee for a process is equivalent to increasing a resource allocation for that process.***" Ex. 3 at 19 (2003-11-17 Response) (emphasis added).[2] The applicant further argued that Vahalia does not anticipate the claims because Vahalia does not disclose increasing a quality of service guarantee. Rather, the applicant argued that "[o]nce Vahalia has scheduled a task, there is no further determination of whether the resource that is handling the task is overloaded. In other words, in Vahalia, the quality

[1] The specifications of the '686, '726 (discussed in Section III), and '937 patents are identical.
[2] *See MBO Laboratories, Inc. v. Becton, Dickinson & Co.*, 474 F.3d 1323, 1327 (Fed. Cir. 2007) (prosecution history of related patents its relevant to the claim construction analysis).

of service of each resource ***is fixed ahead of time and is not altered once a task is scheduled to the resource***." *Id*. (emphasis added).

The PTO relied on the applicant's disavowal regarding this issue, noting that:

> In the claim languages as found in claims 1, 13, 27, 31, 37, 39, and 44, ***Applicant appears to equate 'the quality of [service] guarantee' to 'allocating a portion of resources' and 'increasing the quality of service guarantee' to 'increasing the resources allocated to an associated service.' The slight departure of these phrases from its nominal meanings is being recognized here.*** Although the examiner has adopted Applicant's lexicographer [sic] in the following prior art rejection, it is noted that care must be taken when comparing to its counterpart in the referenced passages (such difference was raised during a recent telephone interview with Applicant's representative).

Ex. 5 (2004-05-10 Office Action) (emphasis added). Thus, in both a response and in a telephonic interview with the examiner, the applicant represented that the terms "resource allocation" and "quality of service guarantee" are equivalent, and the terms "modifying a resource allocation" and "modifying a quality of service guarantee" are equivalent. The examiner unequivocally memorialized this understanding as set forth above, and the applicant did not dispute this interpretation over the subsequent lengthy file history prior to issuance.

IV should be held to this understanding of the claim term. *Digital Biometrics, Inc. v. Identix, Inc.*, 149 F.3d 1335, 1347 (Fed. Cir. 1998) ("The public has a right to rely on such definitive statements made during prosecution. Notice is an important function of the patent prosecution process….").

The claims further support VMware's proposal. For example, claim 1 of the related '726 patent (which is identical to issued claim 1 of the original '937 patent), recites "in response to the virtual server overloaded signal, to modify a resource allocation for the virtual server" and "receive the virtual server resource modification signal and [] determine whether the first physical host is overloaded." In other words, the claims contemplate that after a virtual server is determined to be overloaded, the resource allocation for the virtual server is modified. And *after* the "resource allocation" for the virtual server is already modified, the system may determine that the physical

host is overloaded.[3] Such a step only makes sense if modifying the resource allocation of the virtual server refers to modifying the quality of service guarantee of the virtual server, as contemplated by VMware's proposal. Ex. 2 (Snoeren Decl.) ¶ 31. On the other hand, IV's proposal which construes modifying a resource, such as memory or disk space, would involve modifying the physical host in some way. This does not make sense in context of the claims. *Id*.

The specification further supports VMware's proposal. The specification clearly and consistently specifies that a resource allocation refers to a quality of service guarantee, and modifying a resource allocation means modifying a quality of service guarantee. For example:

> ***A resource allocation for a virtual server is specified as a "quality of service guarantee" for that particular server. Each physical host stores quality of service guarantees for the virtual servers it hosts.*** As a physical host performs processes associated with a particular virtual server, the physical host accesses the stored quality of service information to enable the physical host to request the correct quality of service from the operating system kernel of the physical host.

'686 patent, 4:49-56 (emphasis added). Table 1 of the specification, reproduced below, provides another example of the term "resource allocation" referring to a quality of service guarantee.

TABLE 1

Virtual Server Resource Allocation in FIG. 1

Virtual Server	Resource Allocation
162A	15% of physical host 160A
162B	60% of physical host 160A
162C	10% of physical host 160B
162D	10% of physical host 160B
162E	10% of physical host 160B
162F	20% of physical host 160C
162G	30% of physical host 160C

[3] The parties agree that "the first physical host is overloaded" should be construed as "the first physical host will not support additional resource allocations at that time."

'686 patent, 4:24-35. As shown in this table, each virtual server is configured with a "resource allocation" defined as a percentage of the resources of a physical host—an arrangement that the specification equates to a quality of service guarantee:

> In one embodiment, ***each individual virtual server 162 has a different quality of service guarantee***. ***Different quality of service guarantees are implemented by allocating different amounts of the resources of each physical host machine 160 to servicing each of the virtual servers 162. Physical host 160 resources may be allocated as percentages of the resources of a particular physical host 160***, or as a particular number of units within a physical host 160 (for example, the operating system may be instructed to allocate X cycles per second to process A and Y cycles per second to process B). In the embodiment shown in FIG. 1, physical host 160 resources are allocated to individual virtual servers 162 as percentages of each physical host 160. ***Table 1 lists the resource allocations of each virtual server 162 as shown in FIG. 1***[.]

'686 patent, 4:9-22 (emphasis added). The specification consistently describes the invention in terms of a quality of service guarantee. *See, e.g.,* '686 patent, 2:39-42 ("Thus it is desirable to provide a system and method for a virtual server capable of ***providing quality of service guarantees*** for a customer, ***which is also capable of adjusting the quality of service based upon changing customer demand***.") (emphasis added); *id.*, 3:53-55 ("***The term 'virtual server' as used herein refers to a virtual server capable of receiving a quality of service guarantee from a physical host***.") (emphasis added).

By contrast, for the reasons set forth above, IV's proposal is not supported by the claims or the specification, and it attempts to disregard an unmistakable prosecution history disclaimer. Moreover, IV's proposal does not actually construe the term "modify a resource allocation"—instead it appears to only offer a proposed construction of "modify a resource ~~allocation.~~" Specifically, the parties agreed to a construction of the term "resource"[4] and IV has effectively

[4] The parties have agreed that the term resource should be construed as a "set of functions and features of a physical host, such as disk space, memory, network capacity and processing cycles (CPU resources), used in implementing tasks for each virtual server."

taken this construction and inserted the word "modify" before it for the present term. The fundamental problem with the strategy, however, is that it attempts to read out the actual claim language of "resource allocation." This is improper. *See Pause Tech., LLC v. TiVo, Inc.*, 419 F.3d 1326, 1334 (Fed. Cir. 2005) ("[i]n construing claims, . . . we must give each claim term the respect that it is due"); *Merck & Co. v. Teva Pharm. USA, Inc.*, 395 F.3d 1364, 1372 (Fed. Cir. 2005) ("A claim construction that gives meaning to all the terms of the claim is preferred over one that does not do so."). Resource *allocation* has a specific meaning according to the intrinsic record, and a POSITA would have understood that a *resource* is very different from a *resource allocation*, and that *modifying a resource* is very different from *modifying a resource allocation*. Ex. 2 (Snoeren Decl.) ¶¶ 28-29. For this additional reason, IV's proposal should be rejected.

B. "resource unavailable messages resulting from denied requests to modify a resource allocation" ('686 patent claims 5-7)

VMware Proposal	IV Proposal
"indications that requests by the virtual server for additional resources are either implicitly or explicitly denied, resulting from denied requests to modify a resource allocation" *See also* construction of "modify a resource allocation"	See IV proposals for "resource unavailable messages" and "denied requests to modify a resource allocation" "resource unavailable messages" = "an indication that a request by the virtual server cannot be immediately serviced" "denied requests to modify a resource allocation" = "a request by the virtual server that cannot be immediately serviced"

Both parties have proposed this term for construction. The fundamental dispute regarding this term surrounds IV's attempt to read the phrase "resulting from denied requests to modify a resource allocation" completely out of the proposed term and the claims. As illustrated in the table above, IV's proposal for "resource unavailable messages" is ***nearly identical*** to its proposal for "denied requests to modify a resource allocation." In fact, IV's proposal for the former term renders the latter term obsolete. The Court should reject IV's attempt to remove the language

"resulting from denied requests to modify a resource allocation" from the claims and instead adopt VMware's proposed construction which gives due credit to the each and every word in this proposed claim term. *See, e.g., Pause Tech.,* 419 F.3d at 1334; *Merck*, 395 F.3d at 1372.

As to the claim language "resource unavailable messages," VMware and IV's respective proposed constructions are the same as the respective proposed construction for the disputed '726 patent term "resource denial". *See* section III.B. For the reasons set forth below in section III.B, VMware requests that the Court adopt its proposed construction of "an indication that a request by the virtual server for additional resources is either implicitly or explicitly denied" for both terms.

As to the claim language "resulting from denied requests to modify a resource allocation," a proper construction should give meaning to this phrase, rather than ignore it altogether as IV proposes. Specifically, IV's proposals, when read together, would construe "resource unavailable messages resulting from denied requests to modify a resource allocation" as "an indication that a request by the virtual server cannot be immediately serviced [resulting from] *a request by the virtual server that cannot be immediately serviced*." Not only is this proposed construction circular, but it also reads out the claim language reciting that there is a denied request to ***modify*** **a** resource ***allocation***. A "request by the virtual server that cannot be immediately serviced" (as IV proposes) is clearly not the same thing as denying a "request to modify a resource allocation" (as the claim recites). Ex. 2 (Snoeren Decl.) ¶¶ 32-33. Indeed, IV's proposed construction of this term even ignores its own proposed construction of the term "modify a resource allocation." IV's proposed construction which attempts to remove critical language in the claims should be rejected.

C. "determination that a virtual server is overloaded" ('686 patent claims 5-7)

VMware Proposal	IV Proposal
"determination that an average number of resource denials for a virtual server is beyond a pre-configured threshold"	Plain and ordinary meaning

See also construction of "resource denials"	

This technical term merits construction beyond a "plain and ordinary meaning." The independent claims of the '686 patent (and '726 patent) discuss both virtual server overload and physical host overload. The parties have agreed upon a construction of physical host overload because it is a technical term described in the specification.[5] By this same logic, the Court should construe the technical concept of a virtual server being overloaded as used in the '686 patent.

The concept of overload of a virtual server is described *very differently* in the specification than the concept of overload of a physical host, and the terms therefore should be construed differently. VMware's proposal applies the clear definition of this term from the specification:

> A determination is made 220 as to whether a particular virtual server resource is overloaded. The number of times a particular resource denial is received in a time window is averaged using one of a number of well-known techniques. ***If the average number of denials is beyond a pre-configured threshold, the virtual server is determined 220 to be overloaded for the corresponding resource***. If the virtual server is not determined to be overloaded, the method continues to monitor 210 virtual server resource denials.

'686 patent, 5:42-50 (emphasis added). Adopting any other construction for this term would render the claim invalid for lack of written description. For example, IV's proposed construction of "plain and ordinary meaning" will likely be argued as equivalent to its definition for "virtual server overloaded signal" in the '726 patent as "an indication that a virtual server has been or is being denied resources." But as argued, IV's proposal lacks foundation in the specification, ignores the clear description of this term (reproduced above), and disregards the relationship between virtual server overload and resource denials that is repeatedly set forth in the intrinsic record.

[5] The parties have agreed on a construction of the '686 patent claim term "indication that a first physical host is overloaded" as "indication that a first physical host would not support additional resource allocations at that time" and the '726 patent claim term "the first physical host is overloaded" as "the first physical host will not support additional resource allocations at that time."

IV has refused to explain what it believes the plain and ordinary meaning of this claim term means, or otherwise explain why it disagrees with VMware's proposal. The Court should not construe the term as "plain and ordinary meaning" where, as here, there is a dispute as to the meaning. *O2 Micro*, 521 F.3d at 1351. Therefore, the Court should adopt VMware's proposal.

D. "virtual server" ('686 patent claims 5-7)

VMware Proposal	IV Proposal
"a process executing on a host computer that accepts communication requests, and that is capable of receiving a quality of service guarantee from a physical host"	plain and ordinary meaning; in the alternative: "a virtual machine that resides on a physical server and uses the physical server's resources but has the appearance of being a separate dedicated machine"

The specification defines that: "The term 'virtual server' ***as used herein*** refers to a virtual server capable of receiving a quality of service guarantee from a physical host." '686 patent, 3:53-55. Despite this, IV's proposed construction does not include the language "capable of receiving a quality of service guarantee from a physical host" and should be rejected on that basis alone. The parties also disagree as to the meaning of the term "virtual server" itself. This term is found in the claims of the '686 and '726 patents, which reissued from a patent originally assigned to the company Ensim and filed on May 11, 2000. This term is also found in the claims of the '051 patent, which reissued from a patent originally assigned to Ensim and filed on March 15, 2000. For the reasons set forth below in section V.A, VMware's proposal of the term "virtual server" is necessitated by the intrinsic record and finds additional support in the extrinsic evidence, while IV's proposed construction relies exclusively on extrinsic evidence which it has then modified.

E. "determining that a second physical host can accommodate the requested modified resource allocation" ('686 patent claims 5-7)

VMware Proposal	IV Proposal
Indefinite, or in the alternative: "determining that a second physical host can accommodate the denied requests to modify a resource allocation"	plain and ordinary meaning; in the alternative:

See also construction of "modify a resource allocation"	"determining that a second physical host can accommodate the request(s) by the virtual server that could not be immediately serviced"

This claim term renders the claim 5-7 indefinite. The term includes the phrase "the requested modified resource allocation." However, the claim never introduces "a requested modified resource allocation." *Baldwin Graphic Sys., Inc. v. Siebert, Inc.,* 512 F.3d 1338, 1343 (Fed. Cir. 2008) ("[MPEP] 2173.05(e) describes the need, in most cases, for claim terms to have proper antecedent bases: 'The lack of clarity could arise where a claim refers to "said lever" or "the lever," where the claim contains no earlier recitation or limitation of a lever and where it would be unclear as to what element the limitation was making reference.'"). IV's proposal seeks to construe the term "the requested modified resource allocation" as "the request(s) by the virtual server that could not be immediately serviced," which is nearly identical to its proposals for the terms "resource denials" and "resource unavailable messages." IV's proposed construction has a number of problems. First, it introduces yet another antecedent basis problem by proposing "the request(s)" when the claim language does not recite "a request(s)." Indeed, this apparent change from the singular of the term "the requested modified resource allocation" to IV's plural construction appears to be a concession that the claim term at least lacks antecedent basis as written. Moreover, IV's understanding of this claim term supports the argument that this claim term should, in fact, be construed as indefinite, because it is not possible to assess for the purposes of infringement, which of "the request(s) by the virtual server that could not be immediately serviced" that the second physical is being determined to be able to accommodate. Furthermore, as discussed above, IV's proposal for this term is at odds with its own proposal for the term "modified resource allocation." Although VMware believes this term to be indefinite for the reasons set forth above, it has also proposed an alternative to remedy the antecedent basis problem

by referring back to the term "denied requests to modify a resource allocation" which is introduced in the prior claim limitation.

F. "a component configured to receive an indication that a first physical host is overloaded, wherein the indication is based on a determination that a virtual server is overloaded and wherein the determination that a virtual server is overloaded is based on one or more resource unavailable messages resulting from denied requests to modify a resource allocation" ('686 patent claim 7)

VMware Proposal	IV Proposal
Means-plus function term. This term is indefinite. Function: receive an indication that a first physical host is overloaded, wherein the indication is based on a determination that a virtual server is overloaded and wherein the determination that a virtual server is overloaded is based on one or more resource unavailable messages resulting from denied requests to modify a resource allocation Structure: This term is indefinite for a lack of sufficient corresponding structure in the specification.	Not subject to § 112 ¶ 6 - in the alternative: Function: receive an indication that a first physical host is overloaded Structure: Dynamic Resource Configuration Module 100; Physical Hosts 160A-C; Virtual Servers 162A-G

The Court should construe this language which recites "a component configured to … [perform a function]" as subject to 35 U.S.C. § 112 ¶ 6. The term "component" is a nonce term that fails to carry a structural meaning. Ex.2 (Snoeren Decl.) ¶¶ 35-37; *Alarm.com, Inc. v. SecureNet Techs., LLC*, No. CV 15-807-RGA, 2019 WL 3996883, at *6 (D. Del. Aug. 23, 2019) ("While the claim term does not use the words "means," the word "component" is a "nonce" or non-structural word under § 2181 of the Manual of Patent Examining Procedure."); *Amdocs (Israel) Ltd. v. Openet Telecom, Inc.*, No. 110CV910LMBJFA, 2018 WL 1699429, at *19 (E.D. Va. Apr. 6, 2018) ("'component' does not refer to any specifically known structure in the art […and] is subject to § 112(f)"). Indeed, the term "component" is not found anywhere in the '686 patent aside from claim 7. *See generally*, '686 patent.

Furthermore, claim language reciting what the component is "configured to" do is functional. *MTD Prod. Inc. v. Iancu*, 933 F.3d 1336, 1343 (Fed. Cir. 2019) ("the claim language

reciting what the mechanical control assembly is "configured to" do is functional.")

However, the term renders the claim indefinite because the '686 patent does not describe any component that performs the claimed function of "receiv[ing] an indication that a first physical host is overloaded, wherein the indication is based on a determination that a virtual server is overloaded and wherein the determination that a virtual server is overloaded is based on one or more resource unavailable messages resulting from denied requests to modify a resource allocation." There is absolutely no software algorithm disclosed in the original specification for determining that a virtual server is overloaded ***based on one or more resource unavailable messages resulting from denied requests to modify a resource allocation***. *See generally*, '686 patent; Ex. 2 (Snoeren Decl.) ¶¶ 38-39. Furthermore, IV has declined to identify ***any*** portion of the '686 patent that discloses an algorithm for the claimed component. Nor is IV's broad proposal of ***multiple*** components ("Dynamic Resource Configuration Module 100; Physical Hosts 160A-C; Virtual Servers 162A-G") sufficient for providing the requisite structure. *Augme Techs., Inc. v. Yahoo! Inc.*, 755 F.3d 1326, 1338 (Fed. Cir. 2014) ("Simply disclosing a black box that performs the recited function is not a sufficient explanation of the algorithm required to render the means-plus-function term definite.") Since this term is a means-plus-function term and the '686 patent fails to disclose any algorithm to perform the recited function, this term renders claim 7 indefinite.

G. "a component configured to determine that a second physical host can accommodate the requested modified resource allocation" ('686 patent claim 7)

VMware Proposal	IV Proposal
Means-plus function term. This term is indefinite. Function: determine that a second physical host can accommodate the requested modified resource allocation *See also* alternate construction of this term in Section II.C	Not subject to § 112 ¶ 6 - in the alternative: Function: determine that a second physical host can accommodate the requested modified resource allocation

Structure: This term is indefinite for a lack of sufficient corresponding structure in the specification.	Structure: Dynamic Resource Configuration Module 100; Physical Hosts 160A-C; Virtual Servers 162A-G

For the reasons discussed in section II.F, this term is a means-plus-function term subject to 35 U.S.C. § 112 ¶ 6. However, the '686 patent does not describe any component that performs the claimed function of "determin[ing] that a second physical host can accommodate the requested modified resource allocation." As discussed in section II.C, to the extent this claim limitation is not determined to be indefinite, the term "the requested modified resource allocation" should be construed to be "the denied requests to modify a resource allocation." There is no software algorithm, as depicted by a flowchart (or otherwise), that is disclosed in the specification for determining that a second physical host can accommodate the denied requests to modify a resource allocation. *See* '686 patent; Ex. 2 (Snoeren Decl.) ¶¶ 40-41.

Again, IV has declined to identify portions of the '686 patent that described an algorithm for the claimed component. As such, this term renders claim 7 invalid as indefinite.

H. "a component configured to generate a physical host transfer signal that indicates a second physical host and to transfer the virtual server from the first physical host to the second physical host if the first physical host is overloaded" ('686 patent claim 7)

VMware Proposal	IV Proposal
Means-plus function term. Function: generate a physical host transfer signal that indicates a second physical host and to transfer the virtual server from the first physical host to the second physical host if the first physical host is overloaded *See also* construction of "virtual server" Structure: Dynamic Virtual Server Mover 140 as described in '686 Patent, 12:1-28; Figure 6.	Not subject to § 112 ¶ 6 - in the alternative: Function: generate a physical host transfer signal Structure: Dynamic Resource Configuration Module 100; Physical Hosts 160A-C; Virtual Servers 162A-G

For the reasons discussed in Section II.F, this term is a means-plus-function term subject to 35 U.S.C. § 112 ¶ 6. However, the '686 patent does identify a software algorithm that is associated with performing the claimed function. Figure 6 and the corresponding description at '686 patent, 12:1-28, describe a component called a "Dynamic Virtual Server Mover" "for transferring a virtual server from one physical host to another physical host." '686 patent, 12:1-3.

As shown in the '686 patent's Figure 6 and at 12:1-28, the "Dynamic Virtual Server Mover" is disclosed as a specific algorithm that involves storing state information, stopping virtual server processes in the first host, accessing the state information by the second host, starting virtual server processes on the second host, and transferring a user of the virtual server from the first host to the second host. The '686 patent discloses the mover "may use either 'make, then break' timing or 'break, then make' timing for the transfer process." '686 patent at 12:27-28. Given that this is the only disclosure in the patent—identified by either side—of an algorithm for performing the claimed function, the Court should adopt VMware's proposal. Ex. 2 (Snoeren Decl.) ¶¶ 42-49.

III. DISPUTED TERMS FROM U.S. PATENT NO. RE42,726 (the "'726 patent")

A. Terms that overlap with disputed claim terms in the '686 patent

The terms identified in Ex. 7 overlap with proposed constructions for terms in the '686 patent discussed above. For the reasons set forth in sections II.A, II.C, and II.D, the Court should adopt VMware's proposed constructions for these terms.

B. "resource denials" ('726 patent claims 1, 4-5 and 8)

VMware Proposal	IV Proposal
"indications that requests by the virtual server for additional resources are either implicitly or explicitly denied"	"an indication that a request by the virtual server cannot be immediately serviced"

To streamline the dispute, the parties have reached agreement regarding the first part of the construction of "resource denials" – i.e., "indications that requests by the virtual server…" (In

contrast with IV, VMware's proposal for this first part of the construction is plural in line with the term resource denial***s*** - i.e., requiring more than one resource denial). The parties, however, disagree on the last part of the construction. VMware's and IV's proposals are drawn from different parts of the specification. VMware's proposal comes from a definitive statement in the specification of "resource denials" – "The virtual server resource monitor 110 monitors different types of resource denials, ***which are*** instances wherein a request for ***additional resources is either implicitly or explicitly denied***." '726 patent, 7:51-54 (emphasis added). By contrast, IV's proposal is drawn from a more tentative description of "resource denials" – "A resource denial ***may refer*** to any request by the virtual server ***that cannot be immediately serviced***…." '726 patent, 2:55-61 (emphasis added). As such, the Court should adopt VMware proposed construction of this term.

C. "quality of service guarantee" ('726 patent claims 1 and 4)

VMware Proposal	IV Proposal
"information that specifies a guaranteed amount of an assigned resource, and that can be dynamically increased/modified"	"a guaranteed resource allotment which can be dynamically increased/modified"

The parties seem to agree that a quality of service guarantee refers to information specifying a guaranteed amount of a resource assignment/allotment that can be dynamically increased/modified. A POSITA would have understood that specifying a "guaranteed" amount of a resource refers to providing a particular amount of an assigned resource that is guaranteed. Ex. 2 (Snoeren Decl.) ¶ 51. This specification of a guaranteed amount can be dynamically increased/modified to specify additional assigned resources if needed. *Id*. The specification explains that conventional static resource assignments would cause a customer who "initially request[ed] a very low level of resources" to "quickly bump up against the limit of the server resources he originally requested" and that therefore "it is desirable to provide a system and method for a virtual server capable of providing quality of service guarantees for a customer, which

is also capable of adjusting the quality of service based upon changing customer demand." '726 patent, 2:3-34. The specification further describes quality of service guarantees as follows.

> Different quality of service guarantees are implemented by allocating different amounts of the resources of each physical host machine 160 to servicing each of the virtual servers 162. **Physical host 160 resources may be allocated as percentages of the resources of a particular physical host 160, or as a particular number of units within a physical host 160** […]. In the embodiment shown in FIG. 1, physical host 160 resources are allocated to individual virtual servers 162 as percentages of each physical host 160. Table 1 lists the resource allocations of each virtual server 162 as shown in FIG. 1:

TABLE 1

Virtual Server Resource Allocation in FIG. 1

Virtual Server	Resource Allocation
162A	15% of physical host 160A
162B	60% of physical host 160A
162C	10% of physical host 160B
162D	10% of physical host 160B
162E	10% of physical host 160B
162F	20% of physical host 160C
162G	30% of physical host 160C

'726 patent, 3:67–4:25 (emphasis added).

> As shown in FIG. 1, virtual server 162A consumes 15% of the physical host 160A resources. **This means that 15% of physical host 160A's disk space, memory, network bandwidth, and CPU processing will be dedicated to servicing the needs of virtual server 162A**. […]
>
> **Each physical host stores quality of service guarantees for the virtual servers it hosts.** As a physical host performs processes associated with a particular virtual server, the physical host accesses the stored quality of service information to enable the physical host to request the correct quality of service from the operating system kernel of the physical host.

'726 patent, 4:32–46 (emphasis added). VMware's proposed construction of this term should be adopted because it is consistent with the specification and provides clarity on the term guaranteed. By contrast, IV's proposal provides no clarity on what is being guaranteed, and thus does not address the apparent dispute between the parties on this term.

D. "a virtual server resource monitor [communicatively coupled to the first physical host and] configured to monitor resource denials and to send a virtual server overloaded signal in response to the resource denials" ('726 patent claims 1 and 5) // "program code for creating a virtual server resource monitor communicatively coupled to the first physical host and configured to monitor resource denials and, in response to the resource denials, to send a virtual server overloaded signal" ('726 patent claim 4)

VMware Proposal	IV Proposal
Means-plus function terms. Function: [creating a virtual server resource monitor communicatively coupled to the first physical host and] monitor resource denials and to send a virtual server overloaded signal in response to the resource denials *See also* construction of the terms "resource denials" and "virtual server overloaded signal" Structure: Virtual Server Resource Monitor 110 as described in '726 Patent, 5:21-65; 7:41-9:46; Figure 3.	Not subject to § 112 ¶ 6 - in the alternative Function monitor resource denials and send a virtual server overloaded signal in response Structure Dynamic Resource Configuration Module 100; Physical Hosts 160A-C

The Court should construe the language in claims 1 and 5 which recites "a virtual server resource monitor … configured to … [perform a function]" as subject to 35 U.S.C. § 112, ¶ 6. The term "virtual server resource monitor" is a black-box term that fails to carry a structural meaning to a POSITA. Ex. 2 (Snoeren Decl.) ¶¶ 53-54; *compare Personal Audio, LLC v. Apple, Inc.*, No. 9:09-cv-111, 2011 WL 11757163, at *21 (E.D. Tex. Jan. 31, 2011) ("If 'computer' or 'processor' is insufficient structure to define the scope of a means-plus-function limitation, the word 'processor' cannot describe sufficient structure when recited directly in a claim limitation itself"). Furthermore, as discussed above, claim language reciting what the component is "configured to" do is functional. *MTD Prod. Inc.*, 933 F.3d at 1343.

Similarly, the Court should construe the language in claim 4 which recites "program code for [performing a function]" as subject to 35 U.S.C. § 112 ¶ 6. The term "program code" is a nonce term that fails to carry a structural meaning. Ex. 2 (Snoeren Decl.) ¶ 55; *Glob. Equity Mgmt.*

(SA) Pty. Ltd. v. Expedia, Inc., 2016 WL 7416132, at *29 (E.D. Tex. Dec. 22, 2016) ("the 'program code for configuring ...' term is governed by § 112, ¶ 6."). Indeed, the term "program code" is not found anywhere in the '726 patent aside from claim 4. *See generally*, '726 patent.

The '726 patent identifies a software algorithm that is associated with performing the claimed function. '726 patent, Figure 3, 5:21-65, 7:41-9:46. Specifically the patent describes a "Virtual Server Resource Monitor" which "monitors different types of resource denials" and "determine[s] if any resources from [a] virtual server [] are overloaded." *Id*. at 7:51-53.

The "Virtual Server Resource Monitor" is disclosed with a specific algorithm that involves monitoring different types of resource denials by selectively intercepting system calls, storing resource denial signals in individual resource denial table(s), determining if the number of resource denials in a time window exceeds a pre-specified threshold to determine if a virtual server is overloaded, and if so, sending a virtual server overloaded signal to the virtual server resource modifier. As this is the only disclosure in the patent of an algorithm for performing the claimed function, the Court should adopt VMware's proposal. Ex. 2 (Snoeren Decl.) ¶¶ 56-61.

E. "a virtual server resource modifier [communicatively coupled to the first physical host and] configured to receive the virtual server overloaded signal and, in response to the virtual server overloaded signal, to modify a resource allocation for the virtual server and to send a virtual server resource modification signal" ('726 patent claims 1 and 5); "program code for creating a virtual server resource modifier communicatively coupled to the first physical host and configured to receive the virtual server overloaded signal and, in response to the virtual server overloaded signal, to modify a resource allocation for the virtual server and to send a virtual server resource modification signal" (claim 4)

VMware Proposal	IV Proposal
Means-plus function terms. Function: [creating a virtual server resource modifier communicatively coupled to the first physical host] receive the virtual server overloaded signal and, in response to the virtual server overloaded signal, to modify a resource allocation for the virtual server and to send a virtual server resource modification signal	Not subject to § 112 ¶ 6 - in the alternative Function receive the virtual server overloaded signal and in response, modify a resource allocation for the virtual server and send a

See also constructions of the terms "virtual server overloaded signal," "modify a resource allocation," and "virtual server" Structure: Virtual Server Resource Modifier 120 as described in '726 Patent, 3:43-48, 3:66-5:4; 5:21-65; 9:47-10:52; Figure 4.	virtual server resource modification signal Structure Dynamic Resource Configuration Module 100; Physical Hosts 160A-C; Virtual Servers 162A-G

The Court should construe this language which recites "a virtual server resource modifier … configured to … [perform a function]" as subject to 35 U.S.C. § 112 ¶ 6. The term "virtual server resource modifier" is a black-box term that fails to carry a structural meaning. Ex. 2 (Snoeren Decl.) ¶¶ 62-63. For the reasons discussed in Section III.D, this term is a means-plus-function term subject to 35 U.S.C. § 112 ¶ 6.

Furthermore, for the reasons discussed in Section III.D, the term "program code for [performing a function]" is a means-plus-function term subject to 35 U.S.C. § 112 ¶ 6. *Id*. at ¶ 64.

The '726 patent identifies a software algorithm that is clearly associated with performing the claimed function. Figure 4, 3:43-48, 3:66-5:4, 5:21-65, and 9:47-10:52 of the '726 patent, describe a component called a "Virtual Server Resource Modifier" which "dynamically modifies the resource allocations of the virtual servers [] on an as-needed basis." '726 patent, 5:7-9.

The "Virtual Server Resource Modifier" is disclosed with a specific algorithm that involves receiving a virtual server overloaded signal from the virtual server resource monitor, determining whether the received signal falls within a pre-specified hysteresis time window, and if not, increasing the virtual server resource allocation (i.e., quality of service guarantee). As this is the only disclosure in the patent—identified by either side—of an algorithm for performing the claimed function, the Court should adopt VMware's proposal. Ex. 2 (Snoeren Decl.) ¶¶ 65-69.

F. "a load balanc[ing/er] [module] [communicatively coupled to the plurality of physical hosts and] configured to receive the virtual server resource modification signal and to determine whether the first physical host is overloaded and, in response to a determination that the first physical host is overloaded, to send a physical host transfer signal that indicates a second physical host" ('726 patent claims 1 and 5) // "program code for creating a load balancing module communicatively coupled to the plurality of physical hosts and configured to receive the virtual server resource modification signal and to determine whether the first physical host is overloaded and, in response to a determination that the first physical host is overloaded, to send a physical host transfer signal that indicates a second physical host" ('726 patent claim 4)

VMware Proposal	IV Proposal
Means-plus function terms. Function: [creating a load balancing module communicatively coupled to the plurality of physical hosts and] receive the virtual server resource modification signal and to determine whether the first physical host is overloaded and, in response to a determination that the first physical host is overloaded, to send a physical host transfer signal that indicates a second physical host *See also* construction of the term "the first physical host is overloaded" Structure: Physical Host Load Balancing Module 130 as described in '726 Patent, 5:21-65; 6:4-19; 10:53-11:52; Figure 5.	Not subject to § 112 ¶ 6 - in the alternative Function receive the virtual server resource modification signal and determine whether the first physical host is overloaded and in the case that it is send a physical host transfer signal indicating a second physical host Structure Dynamic Resource Configuration Module 100; Physical Hosts 160A-C; Virtual Servers 162A-G

The Court should construe this language which recites "a load balancing module … configured to … [perform a function]" as subject to 35 U.S.C. § 112 ¶ 6. The term "a load balancing module" is a term that fails to carry a structural meaning. Ex. 2 (Snoeren Decl.) ¶¶ 70-71; *Williamson*, 792 F.3d at 1350 ("It replaces the term 'means' with the term 'module' and recites three functions performed by the 'distributed learning control module.' 'Module' is a well-known nonce word that can operate as a substitute for "means" in the context of § 112, para. 6."). Furthermore, as discussed above, claim language reciting what the component is "configured to" do is functional. *MTD Prod. Inc.*, 933 F.3d at 1343.

Furthermore, for the reasons discussed in Section III.D, the term "program code for [performing a function]" is a means-plus-function term subject to 35 U.S.C. § 112 ¶ 6. Ex. 2 (Snoeren Decl.) ¶ 72.

Here, the '726 patent identifies a software algorithm that is associated with performing the claimed function. Figure 5, 5:21-65, 6:4-19, and 10:53-11:52 of the '726 patent, describe a component called a "Physical Host Load Balancing Module" (or "Physical Host Load Balancer") which "periodically monitors the resource usage of a group of physical hosts [] and transfers virtual servers to different ones of these physical hosts [] in order to balance the resource loads between the physical hosts []." '726 patent, 10:56-60.

The "Physical Host Load Balancing Module" is disclosed with a specific algorithm that involves receiving a signal from the virtual server resource modifier, using a load-balancing calculator to process this signal and a signal regarding the current physical host's resource loads, determining if the physical host is overloaded, and identifying a different available physical host using an easiest fit heuristic algorithm. As this is the only disclosure in the patent—identified by either side—of an algorithm for performing the claimed function, the Court should adopt VMware's proposal. Ex. 2 (Snoeren Decl.) ¶¶ 73-78.

G. "a dynamic virtual server mover [communicatively coupled to the plurality of physical hosts and] configured to receive the physical host transfer signal and, in response to the physical host transfer signal, to transfer the virtual server from the first physical host to the second physical host" ('726 patent claims 1 and 5) // "program code for creating a dynamic virtual server mover communicatively coupled to the plurality of physical hosts and configured to receive the physical host transfer signal and, in response to the physical host transfer signal, to transfer the virtual server from the first physical host to the second physical host" ('726 patent claim 4)

VMware Proposal	IV Proposal
Means-plus function terms. Function: [creating a dynamic virtual server mover communicatively coupled to the plurality of physical hosts	Not subject to § 112 ¶ 6 - in the alternative

and] receive the physical host transfer signal and, in response to the physical host transfer signal, to transfer the virtual server from the first physical host to the second physical host *See also* construction of the term "virtual server" Structure: Dynamic Virtual Server Mover 140 as described in '726 Patent, 11:63-12:23; Figure 6.	Function receive the physical host transfer signal and transfer the virtual server from the first physical host to the second physical host Structure Dynamic Resource Configuration Module 100; Physical Hosts 160A-C

The Court should construe this language which recites "a dynamic virtual server mover … configured to … [perform a function]" as subject to 35 U.S.C. § 112 ¶ 6. The term "dynamic virtual server mover" is a term that fails to carry a structural meaning. Ex. 2 (Snoeren Decl.) ¶¶ 79-80. Indeed, a search of all U.S. Patents shows that the term appears in only 3 patents (the '686, '726, and '937 patents). *See* Ex. 6. Furthermore, as discussed above, claim language reciting what the component is "configured to" do is functional. *MTD Prod. Inc.*, 933 F.3d at 1343.

Furthermore, for the reasons discussed in Section III.D, the term "program code for [performing a function]" is a means-plus-function term subject to 35 U.S.C. § 112 ¶ 6. Ex. 2 (Snoeren Decl.) ¶ 81.

The '726 patent identifies a software algorithm that is clearly associated with performing the claimed function. Figure 6 and the corresponding description at '726 patent, 11:63-12:23, describe a component called a "Dynamic Virtual Server Mover" "for transferring a virtual server from one physical host to another physical host." '726 patent, 11:63-65.

As shown in Figure 6 and described in '726 patent, 11:63-12:23, the "Dynamic Virtual Server Mover" is disclosed with a specific algorithm that involves storing state information, stopping virtual server processes in the first host, accessing the state information by the second host, and then starting virtual server processes on the second host. The '726 patent discloses the mover "may use either 'make, then break' timing or 'break, then make' timing for the transfer

process." '726 patent, 12:21-23. Given that this is the only disclosure in the patent—identified by either side—of an algorithm for performing the claimed function, the Court should adopt VMware's proposal. Ex. 2 (Snoeren Decl.) ¶¶ 82-88.

H. "the dynamic virtual server mover is further configured to direct the first physical host to store, in the file system, a set of system files for the virtual server and to direct the second physical host to access, from the file system, the set of system files for the virtual server, thereby transferring the virtual server from the first physical host to the second physical host" ('726 patent claims 3 and 7)

VMware Proposal	IV Proposal
Means-plus function terms. Function: direct the first physical host to store, in the file system, a set of system files for the virtual server and to direct the second physical host to access, from the file system, the set of system files for the virtual server, thereby transferring the virtual server from the first physical host to the second physical host *See also* construction of the term "virtual server" Structure: Dynamic Virtual Server Mover 140 as described in '726 Patent, 11:63-12:23; Figure 6.	Not subject to § 112 ¶ 6 - in the alternative Function direct the first physical host to store a set of system files for the virtual server in the file system and to direct the second physical host to access the set of system files for the virtual server from the file system and transferring the virtual server Structure Dynamic Resource Configuration Module 100; Physical Hosts 160A-C; Virtual Servers 162A-G

For the reasons discussed in Section III.F, the term "dynamic virtual server mover … configured to … [perform a function]" is a means-plus-function term subject to 35 U.S.C. § 112 ¶ 6. Furthermore, for the reasons discussed in Section III.F, the '726 patent identifies a software algorithm that is associated with performing the claimed function in Figure 6 and 11:63-12:23. Ex. 2 (Snoeren Decl.) ¶ 89.

IV. DISPUTED TERMS FROM U.S. PATENT NO. 7,949,752 (the "'752 patent")

A. "exhausted" ('752 patent claims 1, 9 and 24)

VMware Proposal	IV Proposal
"unavailable for reuse"	"used up to the allotted or pre-determined amount"

The disputed claim term is part of a claim limitation that recites: "wherein an amount of the ***service resource is exhausted*** upon being consumed by the network-based agent." '752 patent at claims 1, 9 and 24 (emphasis added).[6] VMware's proposal is supported by a clear disclaimer in the prosecution history where the applicant distinguished prior art having ***reusable*** resources as "the exact opposite" of the claim language:

> Thus, Chou does not teach "***wherein an amount of the service resource is exhausted upon being consumed by the agent***," as recited by claims 86, 92, and 107 or "***wherein a discrete unit of the service resource is exhausted upon being consumed by the agent***," as recited by claim 126. Rather, in contrast to "[a] . . . resource [that] is ***exhausted*** upon being ***consumed***," as recited, using respective language, by claims 86, 92, 107, and 126, Chou teaches the exact opposite -- i.e. that the "cartridges" in Chou are ***reused***. (Chou: col. 15, line 63 - col. 16, line 26.)

Ex. 14 (2010-08-20 Response) at 10 (emphasis in original). As shown, the applicant did not state that the prior art disclosed something ***unrelated*** to the claim term. Instead, the applicant's argument defines the claim term "exhausted" as "the exact opposite" of resources that can be "reused."

VMware anticipates that IV will argue that the prosecution history shows that the applicant abandoned this disclaimer. Not so. As shown in the summary of the prosecution history provided below, the applicant made this disclaimer and did not abandon the argument.

[6] For context, '752 patent provides an example that a "***service resource 25*** related to telephony services (e.g., voice mail and call placement) ***may comprise units of long-distance calling time which are consumed as an agent 22 places one or more calls*** utilizing such services 24." '752 patent at 11:8-12 (emphasis added).

1. Prosecution History Summary

The applicant first sought to broadly claim "using a service and a service resource." Ex. 9 (2009-03-16 Response) at claim 86. However, the examiner found that U.S. Patent No. 7,043,532 ("Humpleman") anticipated these broad claims. Ex. 10 (2009-05-05 Office Action) at 2–3.

The applicant then narrowed the claims to recite "using a service and a service resource configured to be consumed by the agent", and distinguished the prior art as merely using service resources. Ex. 15 (2009-08-05 Response) at 10 ("Humpleman describes that "[e]ach server device may include ***hardware as a resource*** in the network for providing services to the user" . . . not the claimed "means, including the network-based agent, for using a service and a service resource ***configured to be consumed*** by the agent when performing the operation on behalf of the user") (emphasis in original). However, in the next office action, the examiner maintained the prior art rejection. Ex. 11 (2009-10-30 Office Action), 2-3.

The applicant then narrowed the claims further to add "wherein the service resource is exhausted after it is consumed by the agent." Ex. 12 (2009-11-13 Response) at claim 86. The applicant distinguished the prior art by stating "in Humpleman 'hardware' is used as a resource, and is not 'consumed' or 'exhausted'". *Id.* at 10. The examiner responded by withdrawing the Humpleman objection and imposing a new prior art rejection involving U.S. Patent No. 6,247,056 ("Chou"). Ex. 13 (2010-03-16 Office Action) at 4–7. Chou disclosed using reusable software modules called "cartridges." *Id*. at 5.

The applicant then made minor claim amendments and distinguished the prior art as disclosing ***reusable*** resources, which the applicant called the "***exact opposite***" of the "***exhausted***" claim language. Ex. 14 (2010-08-20 Response) at 10 ("in contrast to '[a] ... resource [that] is exhausted upon being consumed,' as recited, . . . ***Chou teaches the exact opposite*** -- i.e. that the 'cartridges' in Chou are ***reused***") (emphasis in the original).

The examiner responded by calling the applicant's argument "misleading" because the '752 patent application also discloses reusable resources (e.g., memory storage space). Ex. 16 (2010-11-08 Office Action) at 2. However, although the examiner did not necessarily agree with the applicant's interpretation, the examiner withdrew its prior art rejection and imposed a new objection involving U.S. Patent No. 5,740,231 ("Cohn"). Cohn disclosed consuming ***non-reusable*** fees for using a service (e.g., similar to the '752 patent application's description, at 11:8-12, of consuming long distance minutes). Ex. 16 (2010-11-08 Office Action) at 3.

The applicant responded by "travers[ing]" and not "acquiescing" to the rejection. Ex. 17 (2011-01-04 Response) at 10. Instead, the applicant amended the claims and argued that the claims as amended "recite ***further features*** that distinguish over the applied reference." *Id.* (emphasis added). The claims were allowed in the following office action.

2. Argument

The prosecution disclaimer here could not be any clearer. The applicant stated that it is the "exact opposite" of the "exhausted" claim language when resources can be "reused". This argument was made specifically for purposes of distinguishing prior art. The applicant never abandoned this disavowal. The Court should adopt VMware's proposal. *See SanDisk Corp. v. Memorex Prods., Inc*., 415 F.3d 1278, 1286 (Fed. Cir. 2005) ("When the patentee makes clear and unmistakable prosecution arguments limiting the meaning of a claim term in order to overcome a rejection, the courts limit the relevant claim term to exclude the disclaimed matter.").

In a prior litigation involving the '752 patent, the Magistrate Judge provided a report and recommendation construing the "exhausted" to be "used up to the allotted or pre-determined amount," opining that the "prosecution history, when considered as a whole, demonstrates that the patentee did not make a clear and unmistakable disclaimer." Ex. 18 (Magistrate Report) at 33-34. IV had cited to *Ecolab, Inc. v. FMC Corp*., 569 F.3d 1335 (Fed. Cir. 2009) for the proposition that

there is no disavowal when an applicant abandons a limiting argument after it was rejected by an examiner. *Id.* at 32. The Magistrate's finding was in error and *EcoLab* is easily distinguishable.

EcoLab involved a unique situation where the applicant made a ***legally incorrect*** argument. The applicant had argued that claims reciting opened-ended "consists essentially of" language encompassed "a single biocide" and not "mixtures of biocides." *Ecolab*, 569 F.3d at 1343. The examiner corrected the legal mistake, and the applicant responded by cancelling the claims without disagreement. *Id*. Under the *EcoLab* facts, it is unclear whether applicant was making a narrowing statement or if it merely misunderstood the law. The *EcoLab* court concluded that "a reasonable reader of this prosecution history could conclude that [applicant's] initial statements . . . were hyperbolic or erroneous, that the Examiner corrected [applicant's] error . . ., that [applicant] recognized its error and never again repeated or relied upon the erroneous rationale . . . ." *Id*.

The '752 patent's prosecution history, in contrast, presents no uncertainty or abandonment. There is no legal misstatement here—instead, the applicant made a purely ***factual*** statement that the "exhausted" limitation is "the exact opposite of" the prior art's "reusable" resources. Additionally, the '752 patent applicant did not cancel the claims without disagreement when the examiner called its arguments "misleading"—instead, the applicant responded by stating that it "respectfully traverse[s] this rejection," and "[w]ithout acquiescing to the propriety of the rejection . . . [made amendments] to recite ***further features*** that distinguish over the applied reference." Ex. 16 (2010-11-08 Response to Final Office Action) (emphasis added).

The present case is also distinguishable because, although the examiner stated that the applicant's arguments were misleading, the examiner accepted the disavowal by shifting its prior art rejection to a new reference (Cohn) that disclosed non-reusable resources. Furthermore, "[i]t is of no moment whether the examiner agreed with the disclaimer." *See, e.g., Uship Intellectual*

Properties, LLC v. United States, 714 F.3d 1311, 1315-16 (Fed. Cir. 2013) ("The analysis focuses on what the applicant said, not on whether the representation was necessary or persuasive: 'Regardless of the examiner's motives, arguments made during prosecution shed light on what the applicant meant by its various terms.'"). The key issue is whether or not the patentee rescinded the disclaimer.

Under the present facts, the applicant has not clearly rescinded the disclaimer of reusable resources. *See, e.g., Hakim v. Cannon Avent Group PLC*, 479 F.3d 1313, 1317-18 (Fed. Cir. 2007) ("Although a disclaimer made during prosecution can be rescinded, permitting recapture of the disclaimed scope, the prosecution history must be sufficiently clear to inform the examiner that the previous disclaimer, and the prior art that it was made to avoid, may need to be revisited."). It is of no moment that applicant amended the claims after the examiner called its arguments misleading. *See Desper Prods. Inc. v. QSound Labs. Inc.*, 157 F.3d 1325, 1335-36 (Fed. Cir 1998) ("That the prosecution shifted to a different focus does not blunt the impact of those remarks made to overcome the prior rejection."). IV's proposal attempts to recapture claim scope which was disavowed by the original applicant and therefore should be rejected.

B. "consumed" (recited in '752 patent claims 1, 9 and 24)

VMware Proposal	IV Proposal
"used up"	"used"

VMware's proposal for this term is supported by narrowing statements made during prosecution and a definitional statement provided in the patent specification.

During prosecution, the '752 patent applicant first tried to claim mere ***use*** of resources. *See, e.g.,* Ex. 9 (2009-03-16 Response) at claim 86 ("means, including the network-based agent, ***for using*** a service and ***a service resource***"). But later, in response to a prior art rejection, the applicant was forced to narrow the claims to recite that the resource is "configured

to be consumed by the agent." Ex. 15 (2009-08-05 Response) at claim 86. In making this amendment, the applicant distinguished prior art with an identified agent (a GUI) configured to merely ***use*** resources (hardware) for providing a service:

> Humpleman describes that "[e]ach server device may include ***hardware as a resource*** in the network for providing services to the user" (emphasis added). Humpleman does not teach "means, including the network-based agent, for using a service and a service resource ***configured to be consumed by the agent*** when performing the operation on behalf of the user," as recited by claim 86 (emphasis added) or "using a service and a service resource ***configured to be consumed by the agent*** when performing the operation on behalf of the user," as recited by claim 92 (emphasis added).

Ex. 15 (2009-08-05 Response) at 10 (emphasis in original). As evidenced by this history, "consumed" cannot simply mean "use" as IV proposes—the claim language was amended to employ both "using" and "consumed" in the same phrase, intending that the two words have different meaning. The applicant then used this narrowing amendment to argue around prior art that disclosed an agent configured to use a resource.

Indeed, the past tense use of the term "consumed" is substantively different than "consume" as used throughout the patent. "Consumed" indicates that a resource was used completely or "used up" to some allotted amount, whereas "consume" refers to the present tense and ongoing "use" of a resource. For instance, "[e]ntities" may "purchase, ***consume, or otherwise use*** services." *See* '752 Patent at 2:2-3 (emphasis added). On the other hand, a "service wrapper 26 ***may maintain a record of the amount of service resource 25 consumed*** by various agents." *Id.* at 12:9-12 (emphasis added). Here, "consume" clearly refers to consuming resources that are still available, while "consumed" refers to a resource that was used up to a certain amount. *See also id.* at 22:30-33 ("each service permission 64 specifies whether the agent 22 is authorized ***to***

consume a particular service resource 25 and, in some instances, the amount of such service resource 25 ***that is allowably consumed*** by that agent 22.").

Indeed, the '752 patent's specification defines how "consumed" is different than mere "use". Specifically, the '752 patent equates the term "consumed" with the term "used up". *See* '752 Patent at 8:21-23 (resources "may be ***'consumed' or 'used up'*** during the operation of network system 2") (emphasis added). This is a definitional statement defining "consumed" as VMware proposes. *See, e.g., Intellectual Ventures I, LLC v. Lenovo Group Ltd*., 365 F.Supp.3d 200, 206 (D.Mass. 2019) ("A term set off by quotation marks is 'often a strong indication that what follows is a definition.'") (quoting *Sinorgchem Co., Shandong v. Int'l Trade Comm'n*, 511 F.3d 1132, 1136 (Fed. Cir. 2007)). This definitional statement is consistent with common usage of the term "consumed." *See* Ex. 27 (Webster definition of consume) (defining as "to do away with completely," "use up," and "to engage fully"). This definition is further supported by the specification's only other use of the term consumed in quotations—where the specification describes that "[a]t least some of service resources 25 may comprise discrete units which are 'consumed' during utilization of the respective resource by an agent 22 . . . [such as] units of long distance calling time". '752 patent, 11:5-12.

IV's proposal is shown improper when the full claim limitation is rewritten replacing the term "consumed" with "used" and "exhausted" with "used up to the allotted or pre-determined amount" as IV proposes. Claim 1 would become incomprehensible, reciting: "wherein an amount of the service resource is ~~exhausted~~ used up to the allotted or pre-determined amount upon being ~~consumed~~ used by the network-based agent" (i.e., an amount of service resource would not be "used up" simply *upon* being used). This issue becomes even clearer when read in the context of "service resources" as defined in the specification. One example of "services resources" is "disk

space for storing e-mail messages." '752 patent at 11:2-5. Disk space is not "used up to the allotted or pre-determined amount" simply upon being "used."

In the prior litigation, the Magistrate Judge opined that "consumed" means "used" and not "used up." *See*, Ex. 18 (Magistrate Report) at 29–30 (arguing that the '752 patent uses "consumed" synonymously with "used"). The Magistrate's finding was in error. Many of the examples cited by the Magistrate simply refer to embodiments where the term "consumed" is included in a list of alternatives. *See, e.g., id*. at 30 ("computational resources 21 which may be expended, consumed, or used during the operation"). None of the examples overcome the narrowing statement made during prosecution or the specification's definitional statement identified above.

C. "service" ('752 patent claims 1, 3, 9 and 24)

VMware Proposal	IV Proposal
"An application that is used by an agent on behalf of a principal"	"Network functionality available to agent(s)/network-based agent(s)"

The Court should adopt VMware's proposal for this term. The specification makes clear that an "agent uses the service on behalf of a principal." *See, e.g.,* '752 Patent at 3:3-5. In the section titled "Services" the specification further states that the services may "comprise one or more software applications providing various capabilities that are available to a principal." *Id.* at 10:17-19. The next sentence explains that "[e]ach service 24 may be utilized by one or more agents 22 in order to perform their respective tasks." *Id.* at 10:19-21.

IV's proposal, in contrast, does not help clarify the claim scope. The relevant claims already recite execution by a "network-based agent", therefore IV's language of "network functionality" adds nothing. Therefore, VMware's proposed construction should be adopted.

D. Means-Plus-Function Terms

Attached as Ex. 28 is a list of claim terms for which the parties dispute only whether citations to the specification should be included in what is provided to the jury. VMware contends

that compliance with 35 U.S.C. § 112 dictates a construction of means-plus-function claims that includes a cite to the specific portion of the specification where the corresponding structure is disclosed. Section 112 states, in relevant part: an element in a claim for a combination may be expressed as a means or step for performing a specified function without the recital of structure, material, or acts in support thereof, and such claim shall be construed to cover ***the corresponding structure***, material, or acts ***described in the specification*** and equivalents thereof. 35 U.S.C. § 112 (emphasis added). As such, the Court should adopt VMware's proposals for the means-plus-function elements. *See Ariba Inc. v. Emptoris, Inc.*, No. 9:07-cv-90 (E.D. Tex. Aug. 7, 2008) (construing means-plus-function terms with citations to the specification); *see also Personal Audio,* 2011 WL 11757163, at *11 (construing means-plus function terms the same way).

Additionally, a dispute remains as to the structure for the means-plus-function term "means for monitoring an amount of the service resource used by the network-based agent" of claim 4:

VMware Proposal	IV Proposal
Function: monitoring an amount of the service resource used by the network-based agent Structure: monitor as described in '752 patent at 16:50-61	Function: monitoring an amount of the service resource used by the network-based agent Structure: Service Wrapper 26

The parties agree on the function for the means term. The disagreement lies in the proposed structure. VMware's proposal is correct. The specification states that the "[m]onitor 50 generally ***functions to monitor the amount of respective service resources*** 25 expended, ***used***, or otherwise consumed ***by one or more agents*** 22 which have been authorized to access the service 24." *See,* '752 patent at 16:50-61 (emphasis added). This disclosure matches the function proposed by both parties. In contrast, the structure that IV points to, the Service Wrapper 26, is disclosed as comprising the monitor ***in addition to other structure and functionality.*** *See, id.* at 16:28-29 (comprises both a "converter 48" and a "monitor 50." *See, id.* at 16:28-61. The converter is

irrelevant to the function of monitoring service resources, as it is generally used to "convert between a computer language (or instruction set) used within agent server 20 and a computer language (or instruction set) used within the respective service 24." *See, id.* at 16:30-33. Therefore, VMware's proposed construction should be adopted.

V. DISPUTED TERMS FROM U.S. PATENT NO. RE43,051 (the "'051 patent")

A. "virtual server" ('051 patent claims 1, 3 and 6)

VMware Proposal	IV Proposal
"a process executing on a host computer that accept communications requests"	Plain and ordinary meaning, or alternatively: "virtual machine(s) that reside(s) on a physical server and use(s) the physical server's resources but has/have the appearance of being a separate dedicated machine(s)"

The parties dispute whether IV can change the meaning of virtual server to address an entirely different technology (virtual machines) that is not discussed anywhere in the intrinsic record, when the intrinsic record instead consistently describes the virtual servers of the invention as relating to processes, like server applications. VMware's proposal is necessitated by the intrinsic record and finds additional support in the extrinsic evidence. IV's proposal disregards critical intrinsic evidence, and instead relies completely on extrinsic evidence to redefine virtual server to mean a virtual machine, which IV's evidence says is software that mimics the performance of a hardware device. Worse, IV selectively modifies the dictionary definition it relies upon to arrive at this proposed construction. A fundamental dispute between the parties is whether a POSITA would have understood "server" to be a process, or a machine. For the reasons explained below, the intrinsic record is unequivocal that a POSITA would have understood that a "server" in the context of the intrinsic evidence refers to a process, not a machine, and that the term virtual server means "a process executing on a host computer that accept communications requests."

The claim language supports VMware's proposal. Claim 1 describes "a host computer containing a plurality of virtual servers," and claim 3 describes "a host computer containing

multiple virtual servers." A POSITA would have understood the mechanism for placing multiple virtual servers on a single host machine, as described in claims 1 and 3, involves executing processes on the host machine. *See* Ex. 2, ¶ 91; Ex. 19 at 39 (describing the Apache web server, which is comprised of "processes to handle client requests"); Ex. 20 at 45 ("Virtual servers are servers that have different addresses but refer to the same Web server.").

For instance, a POSITA would have been familiar with the use of virtual servers to host multiple websites on a single host machine: "The illusion of more web sites existing than actual web servers is created through the use of ***virtual servers*** (hosts). Web hosting service is based on this technique. Web hosting service uses the possibility to create a set of ***virtual servers*** on the same server." Ex. 21 at 14–15 (emphasis added); Ex. 2, ¶¶ 92–97. An important example is Apache, "the most widely used web server for commercial web sites." Ex. 19 at xviii; *see also* '051 patent at 3 (citing to Ex. 29, BEN LAURIE AND PETER LAURIE, APACHE: THE DEFINITIVE GUIDE (1999)). "Apache is an example of a *preforking* server," which "means that the main server starts a pool of processes to handle client requests, rather than forking a new process for each incoming request." Ex. 19 at 39; Ex. 2, ¶¶ 94–97. Each of a plurality of Apache processes running on a host computer can be bound to a specific IP address, in which case each process is a virtual server. In this configuration, "[a]ll connections to the Apache server on this IP address are handled by the ***virtual server*** for this site, which might be one of many virtual sites being hosted on the same [physical] server." *Id.* at 98 (emphasis added); *see also* Ex. 29 at 163, 177, 242–43, 295 (describing how to configure Apache virtual severs); Ex. 2, ¶¶ 93–95.

The specification reinforces, and in fact requires, an understanding that the term virtual server in claims 1 and 3 refers to processes, like those used in the Apache sever application. Ex. 2, ¶¶ 96–97. The '051 patent explains that "a server application executing on a single physical host

can be programmed to process requests made to multiple network addresses." '051 patent at 2:47–21. The specification discloses that the '051 patent "is related to U.S. patent Ser. No. 09/452,286, entitled 'Providing Quality of Services Guarantees to Virtual Hosts,'" which "is incorporated by reference herein in its entirety." *Id.* at 1:16–21. The patentee further provides that the '286 application describes how to create the private virtual servers disclosed in the '051 patent. *Id.* at 3:64–67. The '286 application ultimately issued as U.S. Patent No. 6,976,258.

The '258 patent only describes how to create and configure processes, like the Apache virtual servers. *See* Ex. 2, ¶¶ 96–97. The '258 patent defines "[a] server" as "a process, executing on a dedicated physical services client [that] services client requests for a single network address (physical host) only." '258 patent at 3:23-25. "A server program executing on the host opens a communication transport channel (socket) and allows receipt of incoming communications targeted for any of the multiple network addresses assigned to the host. . . . However, multiple communication requests made to a plurality of network addresses associated with a single physical host require simultaneous service. For this reason, a virtual host server typically accepts the communications requests itself and creates child processes to service the requests." *Id.* at 1:24–44. In other words, the virtual server described in the '258 patent (and thus incorporated into the '051 patent) refers to a server application that uses child processes to service communication requests on a plurality of network addresses. The '258 patent explains that "[t]wo known methods exist for utilizing child processes to service communication requests." *Id.* at 1:45–2:5. The two methods described are "fork after accept" (generating child processes after a request is received), *id.* at 1:45-55, and, as with Apache's virtual servers, *see* Ex. 19 at 39, "fork before accept" (generating child processes in advance to service requests), '258 patent at 1:56–5; Ex. 2, ¶ 96–97.

Whichever method is used to create them, the '258 patent teaches that these processes can be bound to individual network addresses, *see id.* at 2:6–22, which "allows an ISP to utilize one physical host computer to provide host services to multiple customers." *Id.* at 2:38–40. Moreover,

> [b]y utilizing the fork before accept method or the fork after accept method, the virtual host server can service requests to multiple network addresses or domain names. Thus, the functionality of numerous hosts is provided by a single physical host computer, servicing requests made to a plurality of network addresses and domain names by multiple customers.

Id. at 2:40–52, 1:24-44, 1:45-2:5, 3:23-24, 4:36-49. This background—which the patentee incorporated by reference to explain how to create the virtual servers in the '051 patent—provides clear context from which a POSITA would understand that the virtual server of the '051 patent is "a process executing on a host computer that accepts communications requests."

All of these disclosures contradict IV's contention that virtual server means a virtual machine. *See* Ex. 2, ¶ 98. Despite the robust discussions of processes and how to use them as virtual servers in the intrinsic record, IV has produced only a single dictionary definition (from after the priority date of the patent) to justify redefining virtual server as virtual machine. That same reference defines a virtual machine as "[s]oftware that mimics the performance of a hardware device such as a program that allows applications written for an Intel processor to be run on a Motorola chip." *See id.* at 554. There is no support in the claims or the specification for IV's contention that the virtual servers of the '051 patent are software that mimics the performance of a hardware device. IV's proposed definition contradicts the disclosures in the specification that the virtual servers relate to server applications that use child processes to service communications requests—a distinctly software-based solution, not hardware emulation. *See* Ex. 2, (Snoeren Decl.) ¶¶ 98–99. Given the intrinsic evidence, a POSITA would not have understood the '051 patent to use virtual server to mean a virtual machine like that disclosed in IV's extrinsic evidence. *Id.*

The prosecution history further confirms VMware's proposal. Ex. 2, ¶ 100. The '051 patent was a reissue of U.S. Patent No. 6,948,003. During prosecution of the '003 patent, the applicant responded to multiple rejections based on U.S. Patent No. 6,286,047 ("Ramanathan") (Ex. 25) and U.S. Patent No. 6,247,057 ("Barrera") (Ex. 26). *See, e.g.*, Ex. 22 at 14–15 (2003-10-14 Response). The applicant conceded that Ramanathan discloses virtual servers "in the context of determining whether a host supports virtual servers," and further explained that an example of a virtual server in Ramanathan is a "website hosted by a virtual server [that has] a unique IP address." *Id.* Likewise, the applicant tried to differentiate the virtual servers in Barrera because it "requires that the virtual services each have a unique locator ID." *Id.* The applicant never disputed that the virtual servers in Ramanathan and Barrera were the same as those in the application.

Ramanathan and Barrera both describe virtual servers like the Apache virtual servers described above. Ex. 2, ¶ 100. Ramanathan explains that "the virtual servers model" involves hosting "all of the customer websites . . . using a single host machine," Ex. 25 at 31:48–51, and provides example templates addressing Apache virtual servers. *Id.* at 11:44–13:30. Barrera similarly explains that "current server technology allows administrators to run multiple instances of the same service on a single machine." Ex. 26 at 2:1–12. As an example, Barrera also notes that it is common "for a Web server to support thousands of domains on the same Web service," such that "the physical host server can be said to support multiple 'virtual services' on multiple 'virtual servers.'" *Id.* at 2:13–19. Nothing in either Ramanathan or Barrera indicate that a virtual server is anything other than a process, much less software emulating a hardware device. Ex. 2, ¶ 100.

The intrinsic evidence speaks clearly, and the extrinsic evidence supports what it says—virtual server means a process executing on a host computer that accepts communications requests.

B. "physical interface[s]" ('051 patent claims 1 and 3)

VMware Proposal	IV Proposal

"hardware that provides a point of communication between two or more devices"	Plain and ordinary meaning

VMware contends the term "physical interface" needs to be construed to resolve an apparent dispute between the parties. It appears IV and VMware differ over whether the plain meaning of the word physical interface as it is used in the '051 patent means hardware. IV has refused to confirm this is actually in dispute, and has instead opaquely insisted that the term has a plain and ordinary meaning without providing any clarity as to what it believes that meaning is, or whether that meaning comports with VMware's understanding that the physical interface is hardware. As explained below, the proper construction of physical interface is "hardware that provides a point of communication between two or more devices."

The claim language requires VMware's construction. The term physical interface appears in independent claims 1 and 3 to facilitate communication between two or more devices. *See, e.g.*, '051 patent at claim 1 ("receiving . . . a transmission on a physical interface"; "sending the transmission . . . on the determined physical interface"). Notably, the patentee chose to use the term *physical* interface, not just interface. This choice has meaning. Ex. 2, ¶ 104. The term physical is used throughout the specification to differentiate hardware components from virtualized or logical components. For instance, the specification uses the term physical to differentiate between the physical host computers (or physical servers) and the virtual servers running on them. *See, e.g.*, '051 patent. at 2:47–64 ("a server application executing on a single physical host"), 4:35–38 ("one or more private virtual servers located on a physical host computer"), 4:56–65, 7:1–11, 7:47–48, 8:4–5. Likewise, the specification uses the term physical to separate the logical tunnel constructs from the hardware over which the data is transmitted. *See id.* at 10:17–20 ("Each tunnel is not a separate physical connection; it is a specific encapsulation of data allowing the data to be separated out from other data sent on a physical connection."). A POSITA would have thus understood the

patentee's use of physical mean a hardware interface comprised of circuitry. Ex. 2, ¶¶ 102–06; *see also* Ex. 23, (IEEE definition: "circuitry that interfaces a module's nodes to the input link, output link, and miscellaneous signals").

The specification also provides examples of physical interfaces that confirm VMware's proposal, such as a hardware network interface cards. '051 patent at 8:51–57, 9:4–9, 9:56–60. So too do uncontested examples from the examiner in the prosecution history, like "physical copper connections." Ex. 24 (2010-05-17 Response) at 7. There is no indication in the intrinsic record that the physical interface described in the claims is anything other than hardware that provides a point of communication between two or more devices.

C. physical interfaces and tunnel identifiers in the storing / receiving / determining / sending terms ('051 patent claims 1 and 3)

The terms "physical interface(s)" and "tunnel identifier(s)" are each used six times in each independent claim of the '051 patent. Because the claim language is ambiguous about which of the physical interfaces and tunnel identifiers are being referenced at the various steps, this creates a significant antecedent basis problem that makes the claims impossible to decipher without clarification. VMware's proposal for these terms (*see* Ex. 8) offers a simple way to resolve this problem that is fully supported by the specification—identifying each iteration of "physical interface(s)" and "tunnel identifier(s)" as either corresponding to an incoming (received) transmission or an outgoing (sent) transmission.

Claims 1 and 3 use two distinct sets of physical interfaces and tunnel identifiers. One set is of physical interfaces and tunnel identifiers associated with the receiving step; these are the physical interfaces and tunnel identifiers associated with the incoming transmission. When a transmission is received, the method requires identifying from the customer lookup table (or customer lookup information) the customer forwarding table (or customer forwarding information)

that is associated with the pair of the physical interface and the tunnel identifier of the incoming transmission. The method then requires looking in the identified customer forwarding table (or customer forwarding information) for the destination network address of the incoming transmission to determine the pair of a physical interface and a tunnel identifier to use for sending the transmission. The transmission is then sent using the physical interface and tunnel identifier determined from the customer forwarding table (or customer forwarding information).

The physical interfaces and tunnel identifiers stored in the customer lookup table (or customer lookup information) and used in the receiving step are the *incoming* physical interfaces and the *incoming* tunnel identifiers. Figure 8 shows the customer lookup table of the invention, which includes the "Incoming Physical Interface" and the "Incoming Tunnel ID," and which are together associated with a "Customer ID." The specification further explains that the fields in the customer lookup table include the "incoming physical interface," the "incoming tunnel identifier," the "customer identifier," and that "[t]his customer identifier provides an index to the correct customer forwarding table associated with the physical interface/tunnel identifier pair." '051 patent at 11:64–12:1, 12:59–63 ("The incoming physical interface and tunnel identification information is read 740 from the packet 718, and presented 744 to a customer lookup table.").

Similarly, the physical interfaces and tunnel identifiers stored in the customer forwarding table (or customer forwarding information) and used in the sending step are the *outgoing* physical interfaces and the *outgoing* tunnel identifiers. Figure 9 shows the customer forwarding table of the invention, which includes the "Outgoing Tunnel ID" and the "Outgoing Physical Interface," which are together associated with a "Destination IP Address." The specification teaches that the fields in the customer forwarding table includes the "destination IP address," the "outgoing tunnel identifier," and the "outgoing physical interface," and that "[b]ased upon the destination IP address

of the particular transmission, the proper outgoing tunnel and outgoing physical interface is determined." *Id.* at 12:10–20, 11:30–32, 11:44–31, 12:30–33, 13:5–11 ("From the information contained in customer forwarding table 910, the correct outgoing physical interface and tunnel identifier for packet 718 is identified 754."), 14:56–63. VMware's proposals are correct because they use clear disclosures in the specification that would inform the understanding of a POSITA to correct the antecedent basis issues that would otherwise render the claim language indefinite.

VI. DISPUTED TERMS FROM U.S. PATENT NO. RE44,818 (the "'818 patent")

A. "hierarchical token bucket resource allocation" (recited in '818 patent claims 1, 17, 30, 32 and 42)

VMware Proposal	IV Proposal
the specific class-based scheduling algorithm known in the art as the "hierarchical token bucket"	Plain and ordinary meaning

VMware's construction is the only meaning that any POSITA would ascribe the term. The terms "hierarchical token bucket" and "token" as used in the claims of the '818 patent refer to the very specific method of allocating bandwidth resources referred to in the art as the "hierarchical token bucket," or "HTB." Ex. 2, Snoeren Decl. ¶¶ 107-112. HTB was invented or at least popularized by Martin Devera when he implemented it as a queuing discipline in the Linux operating system. *Id*. at ¶ 107. It is not something that the inventors created. In fact, during the prosecution, the inventors submitted an article[7] describing the well-known HTB algorithm in information disclosure statements. *See* '789 Patent File History, 2008.01.15 Information Disclosure Statement; '818 Patent File History, 2013.08.04 Information Disclosure Statement. This is one of many prior art articles (five of which appear on the face of the patent) that demonstrate that HTB is a "proper noun" or "term of art" that refers specifically to the scheduling mechanism known as "hierarchical token bucket" and not anything else. Ex. 2, ¶ 107.

[7] Ex. 43, ("Valenzuela Article").

Token buckets are based around the concept of exchanging logical "tokens" for the right to send data. Tokens are stored in buckets, which are replenished (and drained) at parameterized rates. Dr. Snoeren explains how the HTB uses a specific hierarchical arrangement of token-bucket rate limiters, each of which are a basic mechanism used for traffic management in computer networking. Ex. 2, ¶ 110. In HTB, traffic is divided into classes, and each class is associated with a token bucket in the hierarchy. *Id*. The tokens may be "borrowed" from buckets higher up in the hierarchy in accordance with a well-understood process. *Id*.

IV's position seems to be that the term does not have a specific meaning, but that it could refer to any hierarchical structure for managing bandwidth. Dr. Snoeren explains HTB was one of several well-known and frequently implemented class-based techniques to manage bandwidth resources, each of which generally had its own proper name and each of which was understood to be different. *Id*. at 108.

Dr. Snoeren confirms that the inventors used the term HTB in this customary way and did not intend or attempt to deviate from commonly understood meanings. *Id*. at ¶ 111. The patent itself uses "HTB" as a proper noun and notes that HTB it is one specific form of scheduling and that there are other, unclaimed methods. For instance, it states that "Hierarchical token bucket can be considered as a class based scheduling mechanism" and states that the "QoS manager" will "queue or forward" traffic "using scheduling and queuing methods **such as** hierarchal token bucket (HTB)." '818 patent at 9:51-10:53 (emphasis added). This description mirrors the description in the Velanzuella reference (cited by the '818 patent) and is consistent with how the term was understood in the art. Ex. 2, at ¶ 111. Moreover, during prosecution, applicants argued that the claims were patentable because the HTB was different than other "hierarchical or tree structure for storing resource reservations." Ex. 31 (2009.08.25 Response) at 12. IV disagrees with VMware's

proposed construction but refused to say what was wrong with VMware's construction. Ex. 44. This does not resolve the dispute. VMware's construction should be entered. *See O2 Micro*, 521 F.3d at 1361; *US Foam Inc. v. On Site Gas Systems, Inc.*, 735 F. Supp. 2d 535, 556 (E.D. Tex. 2010) ("Even though the parties have not identified the substance of their position is or the real dispute, the Court must nonetheless fulfill its duty to determine the proper scope of the claims.").

B. "token" ('818 patent claims 1, 17, 30, 32-33, 37-42)

VMware Proposal	IV Proposal
"token as used in hierarchical token bucket resource allocation"	Plain and ordinary meaning

This Court should adopt VMware's proposed construction for the term "token" because it is consistent with how tokens were used in token-bucket traffic management approaches such as HTB at the time of the invention. Ex. 2, Snoeren Decl. ¶ 110. The intrinsic and extrinsic evidence discussed in section VI.A above applies here. In contrast, IV's proposed construction[8] is contrary to well-established Federal Circuit Precedent. *See O2 Micro*, 521 F.3d at 1360-6.

C. "enforc[e/ing]", "receiv[e/ing]", "classify[ing]", "compar[e/ing]", "forward[ing]", and "buffer[ing]" ('818 patent claims 1, 17, 30, 32, 33, 37, 38, 39, 42)

VMware Proposal	IV Proposal
"enforcing . . . across the physical [storage network] interface of the virtual I/O server" "receiv[e/ing] in the virtual I/O server" "classify[ing] in the virtual I/O server" "compar[e/ing] in the virtual I/O server" "forward[ing] in the virtual I/O server" "buffer[ing] in the virtual I/O server"	Plain and ordinary meaning (for each)

[8] In IV's disclosure of extrinsic evidence," IV cited to a dictionary definition for "token." However, a POSITA would understand that this definition is unrelated to the tokens used in the context of token-bucket traffic management approaches such as HTB. Ex. 2 (Snoeren Decl.) ¶ 110.

VMware's proposals for these terms "stays true to the claim language and most naturally aligns with the patent's description of the invention," and thus is the "correct construction." *Phillips*, 415 F.3d at 1316. The Federal Circuit has approved constructions that incorporate features that are "repeatedly and consistently" described as aspects of the invention. *See, e.g., SkinMedica, Inv. v. Histogen, Inc.*, 727 F.3d 1187, 1203-04 (Fed. Cir. 2013) ("clear, repeated, and consistent statements in the specification" limit claim scope); *Microsoft Corp. v. Multi-Tech Systems, Inc.*, 357 F.3d 1340, 1347-48 (Fed. Circ. 2004) (statements were "not limited to describing a preferred embodiment, but more broadly describe[d] the overall inventions").

Here, the patent discloses that "Fig. 4 is a component diagram showing the two-tier hierarchical components **of the virtual I/O server** QoS." '818 patent, 8:40-44 (emphasis added). And figures 5-10, repeatedly and consistently describe how the "receiv[e/ing]," "classify[ing]," "compar[e/ing]," "forward[ing]," and "buffer[ing]" occur in the virtual I/O server. For example, Figure 5 is a flow chart illustrating that these functions occur in the Virtual I/O server. *See* '818 patent Fig. 5 at 502, 504, 510, 516, 526, 528. Likewise, figures 6-10 also illustrate that these functions occur in the Virtual I/O server. In fact, nowhere else in the '818 patent does it describe these functions occurring anywhere else besides the Virtual I/O server.

The "enforc[e/ing]" term is also repeatedly and consistently described in the specification as being enforced across the physical interface of the Virtual I/O server.[9] Although the term "compar[e/ing]" does not explicitly appear in Figures 5-10, when reading the limitation in light of the claims, figures 5-10 refer to a comparison made between the received storage command,

[9] *See, e.g.,* '818 patent, 2:5-9 ("The hierarchy is based on partitioning of network interfaces and I/O subsystems transaction types, with QoS allocation decisions made on each hierarchy independently. This distributed transaction scheme provides scalable and fine-grain QoS management in virtual I/O servers."); *id.*, 8:5-28 ("QoS is performed on I/O communications from application servers 102 in various hierarchical tiers in virtual I/O server 106.").

input/output communication, or local area network packets and tokens which is shown throughout figures 5-10. *See, e.g.*, Fig. 5 at 510 (sufficient tokens to forward?). Moreover, the '818 patent makes clear that "Virtual I/O Server 106 provides the storage and external networking needs of application servers"'818 patent at 3:15-16. Accordingly, a POSITA would understand that the only relevant place for these functions to exist is in the virtual I/O server itself.

As a preliminary matter, the parties do not dispute the meaning "enforce[e/ing]," "receiv[e/ing]," "classify[ing]," "compar[e/ing]," "forward[ing]," and "buffer[ing]." However, as with other terms, IV has simply disagreed with VMware's proposals without providing a substantive explanation. Here, VMware's proposed constructions describe the features of the invention and limits the scope of the invention accordingly. *See Verizon.*, 503 F.3d at 1308. In contrast, IV's "plain and ordinary meaning" proposal failed to provide a meaning that would "assure that the jury understands that it is not free to consider its own meaning for disputed claim terms," thus, it's proposed construction should be rejected. *Sulzer*, 358 F.3d at 1366.

D. "maintaining a connection over a network fabric" ('818 patent claims 1, 17, 30, 32 and 42)

VMware Proposal	IV Proposal
"maintaining a connection between the physical interface of the application server and the physical interface of the virtual I/O server over a network fabric"	Plain and ordinary meaning

VMware's proposal is consistent with how a POSITA understands the term "network fabric." That is, a POSITA would understand that a physical interface connects to a network fabric. Ex. 2, Snoeren Decl. ¶¶ 113-14 ("In one implementation, the I/O fabric stack and I/O fabric PHY interface can employ the Reliable Connections (RC) supported by the Infiniband standard.") (citing '818 patent at 3:9-11). Indeed, the specification discloses that the application server connects to the network fabric and the virtual I/O server also connects to the network fabric. *See*

'818 patent at Fig. 1; *see also id*. at 3:22-23 ("One or more application servers 102 might be connected to the virtual I/O server 106 over I/O switch fabric 104"). As explained below, the intrinsic evidence supports the understanding of a POSITA that the term "connection" refers to the network fabrics connection with the physical interface of the application and the physical interface of the I/O switch fabric.

The specification specifically discloses that the physical interface of the application server connects to the network fabric. To illustrate, Figure 2 depicts the protocol stack of an application server. As shown in Figure 2 of the '818 patent, the application server's "I/O Fabric PHY interface 202" connects with the "I/O switch fabric 104"; *see also, e.g., id*. at 4:32-33 ("I/O fabric PHY interface 202 generally refers to the hardware interface or interconnection to the I/O switch fabric."). Likewise, the physical interface of the virtual I/O device connects to the network fabric. To illustrate, the specification discloses that "Virtual I/O server 106 connects to the I/O switch fabric 104 through I/O fabric interface 110 such as Infiniband ports." '818 Patent at 3:9-11. A POSITA would have known that Infiniband ports are physical interfaces Ex. 2, Snoeren Decl. ¶¶ 113-14 ("Infiniband is a popular networking standard that would have been well known to one of skill the art, and understood to employ physical network interfaces.").

IV disagrees with VMware's proposed construction but will not provide VMware notice as to why. Instead, it proffers that the term should be afforded its "plain and ordinary meaning." IV's proposed construction should be rejected because "when the parties present a fundamental dispute as to the scope of the asserted claims, the Court, and not the jury, must resolve that dispute." *O2 Micro*, 521, F.3d at 1360; s*ee also Reedhycalog UK, Ltd. v. Baker Hughes Oilfield Operations Inc.*, 2008 WL 2152268, at *1 (E.D. Tex. May 21, 2008) ("a court may not decline to construe a claim term or rely on the term's ordinary meaning where such a construction does not resolve the

parties' claim-scope dispute, allowing the parties to present claim scope arguments to the jury"). VMware's proposed construction should be accepted because the "statements describing the [the network fabric connections with the application server and Virtual I/O server as physical] found throughout the specification were not limited to describing a preferred embodiment, but more broadly describe[d] the overall invention." *Microsoft*, 357 F.3d at 1347.

For the reasons stated above, the Court should accept VMware's proposed construction.

E. "virtual storage network interface layer of an application server" / "virtual network interface layer of an application server"/ "virtual interface layer of an application server" ('818 patent claims 1, 17, 30, 32 and 42)

VMware Proposal	IV Proposal
"a virtual storage network interface to higher layers of the virtual node in an application server" / "virtual network interface layer to higher layers of the virtual node in an application server" / "virtual interface layer to higher layers of the virtual node in an application server"	Plain and ordinary meaning

The '818 patent describes the operation of the interfaces for the virtual nodes: "virtual network interface 220 presents a virtual link layer interface to higher layers of the protocol stack." '818 Patent at 5:16-17. VMware's proposed construction is consistent with the intrinsic record. For example, claims 1, 17, and 30 recite "the virtual storage network interface layer is associated with a virtual storage node identifier," and claims 32 and 42 recite "presenting, at a physical interface, a virtual node identifier to a local area network." From the specification, a POSITA would know that a virtual node resides on the application server. *See* '818 patent at 6:27-29 ("In another implementation, an application server is a virtual machine server, hosting one or more virtual machine monitors."). These virtual nodes located on the application servers must have a virtual network interface in order to communicate with the subsystems via the virtual I/O server. *Id*. at 3:25-28 ("Application servers 102 might include one or more virtual network interface modules to enhance the performance of their virtual access with SAN I/O subsystems 114 and

LAN I/O subsystems 116."); *see also id.* at 3:18-20 ("the virtual I/O server 106 creates virtual device interfaces for application servers 102."); Fig. 2 at 208a and 220. Networking protocols are typically described in terms of a protocol stack (*see, e.g.*, '818 patent at Fig. 2), where each "layer" of the stack provides services to those above it, and higher layers make use of those services through an interface. *See* '818 patent at 5:16-17 ("Virtual network interface 220 presents a virtual link layer interface to higher layers of the protocol stack."). A POSITA would therefore understand that the claimed "virtual [storage network] interface layer[s]" correspond to those depicted in Figure 2, which shows the "protocol stack and software modules of an application servers" '818 patent at 2:32-34. A POSITA further understands that a layer in a protocol stack necessarily provides an interface to the layers above it in the stack. Ex. 45, Definition of Layer, VMware-IV_000032322 ("Referring to the protocol or protocols operating at a particular level within a network architecture. Such an architecture commonly is detailed in a protocol stack . . . the bottom layer, the Physical Layer, deals with physical and mechanical aspects of the **interface** between a device and a transmission medium.") (emphasis added). Indeed, the specification does not describe any other protocol stack or interfaces to layers in protocol stacks to which these claim terms would reasonably correspond. '818 Patent at 5:16-17. IV's proposed construction should be rejected because it is contrary to Federal Circuit precedent when the parties have a claim-scope dispute. *O2 Micro*, 521 F.3d at 1361-62. VMware's proposed construction should be adopted because it is consistent with the language of the claims and intrinsic evidence.

F. "one or more input/output virtualization modules comprising computer-readable instructions operative to cause the one or more processors to" performs functions terms ('818 patent claim 17)

Due to the length of these terms, the parties' proposals are included separately in Ex. 32, instead of in a summary table herein. Each of these claim terms include the word "module" which "is a well-known nonce word that can operate as a substitute for "means." *Williamson*, 792 F.3d

at 1350. The claims expand on this nonce term by reciting "input/output virtualization modules comprising computer-readable instructions operative to cause the one or more processors to" perform certain functions, but this additional language similarly "fails to impart any structural significance to the term." *Id*. at 1351. Lacking structural language, the claims should be construed as means-plus-function terms.

However, these claims reciting these means-plus-function terms are invalid as indefinite because the specification also fails to disclose a corresponding structure for performing the recited functions. *See Williamson*, 792 F.3d at 1352. Here, the specification is devoid of any algorithm in the specification or prosecution history for performing the claimed functions. Ex. 2, Snoeren Decl. ¶ 115; *See Grecia v. Samsung Elec. Am., Inc.*, 780 Fed. Appx. 912, 916 (Fed. Cir. Aug. 20, 2019) ("For a computer-implemented means-plus-function term, the corresponding structure is typically the algorithm disclosed in the specification for performing the claimed function"). Additionally, IV has failed to identify any specific structure in the specification and instead broadly cites to multiple components without tying these components to a description of how they perform the function. This is improper. *Augme Techs., Inc. v. Yahoo! Inc.*, 755 F.3d 1326, 1338 (Fed. Cir. 2014).

Since the term "input/output virtualization modules" is not a term that refers to a structure and the '818 patent fails to recite an algorithm to perform the recited function, claim 17 of the '818 patent is invalid as indefinite.

VII. CONCLUSION

For the reasons stated herein, VMware respectfully requests the Court adopt its proposed constructions for the disputed terms and phrases.

Dated: March 6, 2020

Respectfully submitted,

WINSTON & STRAWN LLP

/s/ *Katherine Vidal*

Katherine Vidal
Admitted *Pro Hac Vice*
KVidal@winston.com
Michael R. Rueckheim
Texas State Bar No. 24081129
MRuekheim@winston.com
WINSTON & STRAWN LLP
275 Middlefield Road, Suite 205
Menlo Park, CA 94025
Telephone: (650) 858-6500
Facsimile: (650) 858-6559

Thomas M. Melsheimer
Texas State Bar No. 13922550
M. Brett Johnson
Texas State Bar No. 00790975
MBJohnson@winston.com
Michael A. Bittner
MBittner@winston.com
Admitted *Pro Hac Vice*
Winston & Strawn LLP
2121 N. Pearl St., 9th Floor
Dallas, TX 75201
Telephone: (214) 453-6500
Facsimile: (214) 453-6400

Vivek V. Krishnan
Admitted *Pro Hac Vice*
VKrishnan@winston.com
DaWanna McCray
Admitted *Pro Hac Vice*
DMcCray@winston.com
Winston & Strawn LLP
35 W. Wacker Drive
Chicago, IL 60601
Telephone: (312) 558-5600
Facsimile: (312) 558-5700

William M. Logan
Texas State Bar No. 24106214
WLogan@winston.com
Winston & Strawn LLP
1111 Louisiana Street
Houston, TX 77002
Telephone: (713) 651-2766
Facsimile: (713) 651-2700

ATTORNEYS FOR DEFENDANT VMWARE, INC.

CERTIFICATE OF SERVICE

I hereby certify that, on March 6, 2020, the foregoing document was electronically filed with the Clerk of Court using the Court's CM/ECF system which will send notification of such filing to all counsel of record, including counsel of record for Plaintiffs Intellectual Ventures I LLC and Intellectual Ventures II LLC.

/s/ *Katherine Vidal*
Katherine Vidal

IN THE UNITED STATES DISTRICT COURT
WESTERN DISTRICT OF TEXAS
AUSTIN DIVISION

Intellectual Ventures I LLC and Intellectual Ventures II LLC,	
Plaintiffs/Counter-Defendants,	**Civil Action No. 1:19-CV-01075-ADA**
v.	**JURY TRIAL DEMANDED**
VMware, Inc.,	
Defendant/Counter-Plaintiff.	

DEFENDANT VMWARE, INC.'S RESPONSIVE CLAIM CONSTRUCTION BRIEF

TABLE OF CONTENTS

TABLE OF AUTHORITIES

Page(s)

Cases

Statutes

TABLE OF ABBREVIATIONS

Abbreviation	Full Name
PTO	United States Patent and Trademark Office
POSITA	Person of Ordinary Skill in the Art
Snoeren Decl.	Declaration of Alex Snoeren, Ph.D. Regarding Claim Construction
'686 patent	U.S. Patent No. RE 44,686
'726 patent	U.S. Patent No. RE 42,726
'937 patent	U.S. Patent No. 6,985,937
'937 FH	File History of U.S. Patent No. 6,985,937
'752 patent	U.S. Patent No. 7,949,752
'051 patent	U.S. Patent No. RE 43,051
'818 patent	U.S. Patent No. RE 44,818

TABLE OF EXHIBITS
FILED WITH VMWARE'S OPENING CLAIM CONSTRUCTION BRIEF

Exhibit	Title	Dkt. No.
Ex. 1	Parties' Agreed Constructions to Claim Terms	54-1
Ex. 2	Declaration of Alex Snoeren, Ph.D. Regarding Claim Construction	54-2
Ex. 3	'937 Patent File History, Applicant Arguments dated November 17, 2003	54-3
Ex. 4	U.S. Patent No. 5,933,603 to Vahalia et al.	54-4
Ex. 5	'937 Patent File History, Non-Final Office Action dated May 10, 2004	54-5
Ex. 6	USPTO Patent Search of claim term "dynamic virtual server mover"	54-6
Ex. 7	Terms in the '726 patent that overlap with disputed claim terms in the '686 patent	54-7
Ex. 8	The parties' proposals for the terms in the '051 involving multiple recitations of "physical interface(s)" and "tunnel identifier(s)"	54-8
Ex. 9	'752 Patent File History, Response to Office Action dated March 16, 2009	54-9
Ex. 10	'752 Patent File History, Office Action dated May 5, 2009	54-10
Ex. 11	'752 Patent File History, Office Action dated October 30, 2009	54-11
Ex. 12	'752 Patent File History, Response to Office Action dated November 13, 2009	54-12
Ex. 13	'752 Patent File History, Office Action dated March 16, 2010	54-13
Ex. 14	'752 Patent File History, Response to Office Action dated August 20, 2010	54-14
Ex. 15	'752 Patent File History, Response to Office Action dated August 5, 2009	54-15
Ex. 16	'752 Patent File History, Final Office Action dated November 8, 2010	54-16
Ex. 17	'752 Patent File History, Response to Final Office Action dated January 4, 2011	54-17
Ex. 18	Report and Recommendation of United States Magistrate Judge, Intellectual Ventures v. HCC Ins. Holdings, Inc., Case No. 6:15-cv-660 (E.D. Tex. August 26, 2016) ("'752 Patent Magistrate Report")	54-18
Ex. 19	Charles Aulds, Linux Apache Web Server Administration, 39 (2001) ("Aulds")	54-19
Ex. 20	Barry Nusbaum, WebSphere Application Servers: Standard and Advanced Features 45 (1999)	54-20
Ex. 21	Ludmila Cherkasova, FLEX: Design and Management Strategy for Scalable Web Hosting Service, 14–15 (Oct. 1999)	54-21
Ex. 22	October 14, 2003 Amendment & Remarks, U.S. Patent Appl. No. 09/526,980	54-22
Ex. 23	Physical Interface, The IEEE Standard Dictionary of Electrical and Electronics Terms (6th ed., 1996)	54-23
Ex. 24	RE43,051 (U.S. Patent Appl. No. 11/858,091) Patent File History, May 17, 2010 Office Action	54-24
Ex. 25	U.S. Patent No. 6,286,047 ("Ramanathan")	54-25

Exhibit	Title	Dkt. No.
Ex. 26	U.S. Patent No. 6,247,057 ("Barrera")	54-26
Ex. 27	Webster definition of consume	54-27
Ex. 28	'752 Patent List of Disputed Claim Terms	54-28
Ex. 29	BEN LAURIE AND PETER LAURIE, APACHE: THE DEFINITIVE GUIDE 163, 177, 242–43, 295 (1999)	54-29
Ex. 30	*Virtual Server*, MICROSOFT COMPUTER DICTIONARY (5th ed. 2002), p. 555 (IV-VMWARE-00004290)	54-30
Ex. 31	'789 Patent (Orig. Patent) prosecution history, 2009.08.25 Resp. to Office Action at 12	54-31
Ex. 32	Competing Parties' Proposals for the '818 Patent	54-32
Ex. 33	Grotto Networking, available at https://www.grotto-networking.com/BBQoS.html	54-33
Ex. 34	http://tldp.org/HOWTO/Traffic-Control-HOWTO/index.html	54-34
Ex. 35	https://lartc.org/howto/lartc.qdisc.classful.html#AEN1071	54-35
Ex. 36	Linux Advanced Routing & Traffic Control HOWTO, located at http://www.oamk.fi/~jukkao/lartc.pdf	54-36
Ex. 37	Traffic Control HOWTO, Version 1.0.2, Martin A. Brown, located at https://www.tldp.org/HOWTO/html_single/Traffic-Control-HOWTO/	54-37
Ex. 38	Traffic Control using tcng and HTB HOWTO, Version 1.0.1, Martin A. Brown, April 2006, located at http://linux-ip.net/articles/Traffic-Control-tcng-HTB-HOWTO.html	54-38
Ex. 39	Bavier, et al., Operating System Support for Planetary-Scale Network Services, Proceedings of the First Symposium on Networked Systems Design and Implementation (March 2004), located at https://www.usenix.org/legacy/events/nsdi04/tech/full_papers/bavier/bavier.pdf	54-39
Ex. 40	Benita, Kernel Korner - Analysis of the HTB Queuing Discipline Software, available at https://www.linuxjournal.com/article/7562 as of January 26, 2005 and printed in the Linux Journal, March 2005.	54-40
Ex. 41	U.S. Patent 7,161,904 titled: "System and method for hierarchical metering in a virtual router based network switch" to Hussein et al.	54-41
Ex. 42	Bavier et al, Container-based Operating SystemVirtualization: AScalable,High-performance Alternative to Hypervisors, Conference Paper in ACM SIGOPS Operating Systems Review, January 2007, located at http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.1018.1012&rep=rep1&type=pdf	54-42
Ex. 43	Valenzuela, J.L., et al., "A Hierarchical Token Bucket Algorithm to Enhance QoS in IEEE 802.11: Proposal, Implementation and Evaluation, IEEE, vol. 4, Sep. 2004 ("Valenzuela Article")	54-43
Ex. 44	Email from J. Deblois to M. Rueckheim dated March 3, 2020	54-44
Ex. 45	Webster's New Work Telecom Dictionary, Definition of Layer	54-45

TABLE OF EXHIBITS
FILED WITH VMWARE'S RESPONSIVE CLAIM CONSTRUCTION BRIEF

Exhibit	Title
Ex. 46	U.S. Patent No. 6,351,775 ("Yu")
Ex. 47	*Token*, Microsoft Computer Dictionary (5th ed. 2002), p. 532 (IV-VMWARE-00004257)
Ex. 48	U.S. Patent No. 6,976,258 ("Goyal")
Ex. 49	'937 Patent File History, (January 8, 2004 Office Action)

I. INTRODUCTION

For the reasons set forth herein and in VMware's opening brief, VMware requests that the Court adopt VMware's proposed claim constructions for the disputed terms.

II. DISPUTED TERMS FROM U.S. PATENT NO. RE44,686 (THE "'686 PATENT")

A. "modif[y/ied] [a] resource allocation" / "modify[ing] [the] computer resources allocated to a virtual server" ('686 patent claims 5–7)

VMware Proposal	IV Proposal
"modif[y/ied] [a] quality of service guarantee" / "modify[ing] [the] quality of service guarantee of a virtual server" *See also* construction of "quality of service guarantee"	"modif[y/ied] set of functions and features of a physical host(s) used in implementing tasks for each virtual server" / "modify[ing] a set of the functions and features of a physical host(s) used in implementing tasks for each virtual server"

IV's proposal should be rejected because it is contrary to the intrinsic record, including a clear prosecution history disclaimer, and because it is based on a truncated reading of the disputed claim term. Specifically, IV concedes that it is asking this Court to construe the term "**resource allocation**" in accordance with how the specification refers to the term "**resource**." Dkt. No. 53 at 14[1] ("IV's proposal takes its support directly from the intrinsic record which clearly and unambiguously defines 'resource'."). This is improper. A "resource allocation" is not the same thing as a "resource," and IV has provided no explanation to the contrary. *See also* Dkt. No. 54-2 (Snoeren Decl.) at 10 (explaining that a POSITA would have understood these terms to be different). IV's proposal improperly disregards claim language, contrary to Federal Circuit law. *Merck & Co. v. Teva Pharm. USA, Inc.*, 395 F.3d 1364, 1372 (Fed. Cir. 2005) ("A claim construction that gives meaning to all the terms of the claim is preferred over one that does not do so.").

[1] Page citations herein refer to the ECF page numbering unless otherwise noted.

IV's arguments against VMware's proposal are not persuasive. First, IV argues that VMware's proposal "reads out 'resource allocation' entirely and replaces it with 'quality of service guarantee." Dkt. No. 53 at 15. But, in contrast to IV's proposal, which in fact reads out the term "allocation," VMware proposes construing (not reading out) the term based on a clear prosecution history disclaimer by the patentee that equated "resource allocation" with "quality of service guarantee" and "modifying a resource allocation" with "modifying a quality of service guarantee." *See* Dkt. No. 54 at 12–13. This understanding of the claims was expressly acknowledged by the PTO examiner. *Id*. The specification also includes a definitional statement as well as additional clear and consistent statements equating these terms. *Id.* at 13–16. IV's brief fails to address any of this important evidence.

IV also argues against VMware's proposal because the term "quality of service guarantee" is present in the non-limiting preambles of two of the seven claims at issue. Dkt. No. 53 at 15. To the extent that IV is making a claim differentiation argument, that doctrine is inapplicable here. As an initial matter, claim differentiation is merely a presumption, and cannot overcome the clear disclaimer and definitional statements identified by VMware. *Poly-America, L.P. v. API Industries, Inc.*, 839 F.3d 1131, 1137 (Fed. Cir. 2016). Moreover, the procedural history of these reissue patents provides further reason not to apply the presumption. The two independent claims which include a preamble reciting "the computer resources allocated to the virtual server being specified as a quality of service guarantee" ('726 patent claims 1 and 4) are unchanged from the claims as originally issued in the parent '937 patent, which was prosecuted on behalf of the original assignee Ensim. However, the other five independent claims ('726 patent claims 5 and 8; '686 patent claims 5–7) were newly added in the reissue proceeding prosecuted on behalf of Digital Asset Enterprises, an apparent affiliate of IV. Under this procedural posture, it cannot be presumed

that the invention disclosure envisions separate meanings for these terms. Instead, the history suggests that the reissue revisions were an improper attempt to recapture subject matter that had been disclaimed by the original applicant during prosecution.

B. "resource unavailable messages resulting from denied requests to modify a resource allocation" ('686 patent claims 5–7)

VMware Proposal[2]	IV Proposal
"indications that requests by the virtual server for additional resources are either implicitly or explicitly denied, resulting from denied requests to modify a resource allocation" *See also* construction of "modify a resource allocation"	See IV proposals for "resource unavailable messages" and "denied requests to modify a resource allocation" "resource unavailable messages" = "an indication that a request by the virtual server cannot be immediately serviced" "denied requests to modify a resource allocation" = "a request by the virtual server that cannot be immediately serviced"

IV contends that its construction "stays true to the specific sequence of the claim element as a whole." Dkt. No. 53 at 16. As detailed in VMware's opening brief, this is decidedly not the case. Dkt. No. 54 at 16–17.

Instead, IV's proposed construction for "denied requests to modify a resource allocation" is an attempt to re-write the claim as "denied requests ~~to modify a resource allocation~~." IV concedes this point by asking the Court to construe both "resource unavailable messages" and "denied requests to modify a resource allocation" in accordance with "a general definition of 'resource denials'"—a term that is not tied to the claim language of "requests to modify a resource allocation." Dkt. No. 53 at 16–17. IV's proposal should be rejected on at least this basis alone.

[2] As a preliminary issue, IV's opening brief attributes an incorrect proposed construction to VMware for this term. VMware's actual proposed construction is set forth in this table as well as in VMware's opening brief. Dkt. No. 54 at 16–17.

There are other problems with IV's proposal. First, the claim recites "denied requests to **modify a resource allocation**." Both sides have proposed a construction for "modify a resource allocation" (see section II.A). Tellingly, IV's proposal for the present claim language incorporates neither construction; while VMware's proposed construction stays true to the claim language by including the reference to "resulting from denied requests to modify a resource allocation."

Second, IV's proposal for "denied requests to modify a resource allocation" as "a request by the virtual server that cannot be immediately serviced" completely writes out several words in the claim term itself – including "denied," "modify," and "resource allocation." Specifically, IV's proposal does not in any way construe the term "modify." And it does not in any way construe "resource allocation" (or "resource" or "allocation" for that matter). IV's proposal simply *ignores* the presence of these terms in the claim language. Equally problematic, as noted above, IV offers no explanation for why it is proposing near identical constructions for "resource unavailable messages" and "denied requests to modify a resource allocation."

IV's proposal should therefore be rejected and VMware's proposal, which gives due credit to the each and every word in this claim term, should be adopted.

C. "determination that a virtual server is overloaded" ('686 patent claims 5–7)

VMware Proposal	IV Proposal
"determination that an average number of resource denials for a virtual server is beyond a pre-configured threshold" *See also* construction of "resource denials"	Plain and ordinary meaning

IV contends that this term should be given its plain and ordinary meaning because the term and its constituent parts, including the term "overloaded," would have been understood by a POSITA at the time of the invention. Dkt. No. 53 at 23–24. IV's stance on this term is directly contradicted by the fact that IV has proposed a construction different than "plain and ordinary

meaning" for the similar term "virtual server overload signal" claimed in the '726 patent (which includes in common with the present term, the terms "virtual server" and "overload"). Indeed, IV proposed the term "virtual server overload signal" for construction in the first place. IV's stance is also contradicted by the fact that IV proposed the '686 patent claim term – "indication that a first physical host is overloaded" – for construction, and again, proposed a construction different than "plain and ordinary meaning."[3] For the same reasons that IV believed these other terms require construction, the present term merits a construction as well.

Furthermore, as detailed in VMware's opening brief, overload in the context of a physical host is described very differently from overload in the context of a virtual server in the specification. Dkt. No. 54 at 18–19. As such, a POSITA at the time of the invention would have understood that these terms are different. IV argues that VMware's proposal significantly limits the disputed term by importing limitations from an exemplary embodiment. This is not the case. VMware's proposal is drawn from a clear definition of this term in the specification that does not use the word "embodiment" in the sentence (or the paragraph) in which the definition is found:

> A determination is made 220 as to whether a particular virtual server resource is overloaded. The number of times a particular resource denial is received in a time window is averaged using one of a number of well-known techniques. **If the average number of denials is beyond a pre-configured threshold, the virtual server is determined 220 to be overloaded for the corresponding resource**. If the virtual server is not determined to be overloaded, the method continues to monitor 210 virtual server resource denials.

'686 patent, 5:42–50 (emphasis added). IV's opening brief does not consider this portion of the specification and instead cites to unrelated sections of the specification to suggest, incorrectly, that VMware's proposal is importing an illustrative embodiment in the specification into the claims.

[3] The parties have agreed on a construction of the '686 patent claim term "indication that a first physical host is overloaded" as "indication that a first physical host would not support additional resource allocations at that time" and the '726 patent claim term "the first physical host is overloaded" as "the first physical host will not support additional resource allocations at that time."

It bears noting that for the similar term "virtual server overload signal" found in the '726 patent, IV does not propose a plain and ordinary meaning construction, but instead proposes that this term should be construed as "an indication that a virtual server has been or is being denied resources." However, the specification is devoid of any support for this construction, either in a particular embodiment or otherwise. And IV specifically fails to cite any intrinsic support for its proposed construction. By contrast, VMware's proposed construction is not only supported by the specification, but it is consistent across these respective claim terms in the '686 and '726 patents.

D. "virtual server" ('686 patent claims 5–7)

VMware Proposal[4]	IV Proposal
"a process executing on a host computer that accepts communication requests, and that is capable of receiving a quality of service guarantee from a physical host"	plain and ordinary meaning; in the alternative: "a virtual machine that resides on a physical server and uses the physical server's resources but has the appearance of being a separate dedicated machine"

As detailed below in section V.A, VMware's proposal for this term is consistent with a POSITA's understanding of this term (who would have factored in the clear definition from the specification). *See, e.g.,* '686 patent, 3:53–55 ("term 'virtual server' as used herein refers to a virtual server capable of receiving a quality of service guarantee from a physical host"); Dkt. No. 54-2 (Snoeren Decl.) at 13–14, 42–47.

IV's plain and ordinary meaning proposal should be rejected because the parties clearly dispute what the plain and ordinary meaning is. *O2 Micro Int'l v. Beyond Innovation Techn. Co.*, 521 F.3d 1351, 1362 (2008).

[4] As a preliminary issue, IV's opening brief attributes an incorrect proposed construction to VMware for this term. VMware's actual proposed construction is set forth in this table as well as in VMware's opening brief. Dkt. No. 54 at 19.

IV's alternative proposal should be rejected at least because it ignores the specification's clear definition for this term set forth above. It should also be rejected because it inserts the language "virtual machine" into this claim. The term "virtual machine" is not used anywhere in the intrinsic record of these patents, nor is the phrase "appearance of being a separate dedicated machine." Moreover, as discussed further in section V.A, portions of IV's proposed construction are cherry-picked from dictionary support, while other portions lack any support. IV's use of "virtual machine" in its construction is also prejudicial to VMware (whose name stands for "Virtual Machine"-ware). An improper and unsupported construction like this could cause a jury to prematurely judge the merits of IV's infringement allegations despite VMware's defenses.

The intrinsic record also supports VMware's proposal. For example, IV argues that the '937 patent's file history "similarly makes clear that the term 'virtual server' is being used in its customary manner." Dkt. No. 53 at 18. VMware agrees.

Specifically, IV's argument focuses on statements in the '937 patent's file history regarding a prior art reference – U.S. Patent No. 6,351,775 to Yu ("Yu") (Ex. 46). IV argues that the use of the term "virtual servers" by the original applicant in describing Yu amounts to an acknowledgement "that virtual servers are known in the art" and that the term was being used "consistent[ly] with its known meaning, and not some unconventional or heretofore unknown construct [] coining a 'virtual server.'" Dkt. No. 53 at 18. This is an important concession by IV because Yu describes virtual servers in the same manner detailed below in section V.A – i.e. as process-based server applications, such as web servers which can provide for virtual hosting (*i.e.*, servicing requests for multiple network addresses on a single physical host).

For example, Yu's field of invention explains that "[t]he present invention relates generally to providing load balancing across a collection (or cluster) of servers such as proxy servers and

Web servers in the Internet environment." Ex. 46 at 1:21–24 (emphasis added). Yu further explains that "[t]he number of virtual servers is greater than the actual number of servers in the server cluster. The [Domain Name Server] DNS (167) and [Transmission Control Protocol] TCP router can then **dynamically map each virtual server to one of the actual servers in the cluster**." *Id*. at 11:61–64 (emphasis added).

As such, Yu describes virtual servers as a type of **process**, which unlike virtual machines, can be mapped to a physical host. Dkt. No. 54-2 (Snoeren Decl.) at 42–44. The examiner, who is presumed to make informed findings as to the meaning of prior art references to a POSITA,[5] understood Yu the same way, and explicitly equated a virtual server with a **process**:

> - [It is noted that, as taught by Yu at col.11, lines 54-64, each virtual server (which is a process and can be represented by the identifier (e.g., URL) of the object requested) is assigned a class and through the class-to-server mapping, each virtual server is eventually mapped to an actual server residing in a physical host. In other words, there is mapping between a requested object and a physical host (see col.15, lines 21-37)].

Ex. 49 (2004-01-08 Office Action) at 3; *see also id.* at Dkt. No. 54-5 at 5–6 (2004-05-04 Office Action stating same). Indeed, the original patent applicant never disputed this understanding that a virtual server is a type of process.

E. "determining that a second physical host can accommodate the requested modified resource allocation" ('686 patent claims 5–7)

VMware Proposal	IV Proposal
Indefinite	plain and ordinary meaning; in the alternative:

[5] *In re Berg*, 320 F.3d 1310, 1315 (Fed. Cir. 2003) ("As persons of scientific competence in the fields in which they work, examiners […] are responsible for making findings, informed by their scientific knowledge, as to the meaning of prior art references to persons of ordinary skill in the art and the motivation those references would provide to such persons.")

	"determining that a second physical host can accommodate the request(s) by the virtual server that could not be immediately serviced"

IV's opening brief highlights why the antecedent basis problem in this claim language renders the term indefinite. IV argues that "'requested modified resource allocation' ultimately refers back to the denied request to modify a resource allocation." Dkt. No. 53 at 22. The relevant claim language is copied below:

> wherein the determination that a virtual server is overloaded is based on one or more resource unavailable messages resulting from **denied requests to modify a resource allocation**;
>
> [a component configured to] determin[e/ing] that a second physical host can accommodate **the requested modified resource allocation**;

'686 patent, 14:39–45 (emphasis added). As shown, the first part of the claim refers to "denied requests" (i.e., **more than one** denied request), while the second part of the claim refers to a determination of whether a host can accommodate "the requested modified resource allocation" (i.e., **one** request). As such, even accepting IV's explanation of the antecedent basis, ambiguity remains because the claim offers no guidance as to *which one* of the plurality of denied requests to modify a resource allocation that "the requested modified resource allocation" refers back to. Nor does IV cite anything in the specification that would offer guidance to a POSITA on this issue.

Furthermore, IV's alternative proposal that seeks to construe the term "the requested modified resource allocation" the same way as it proposed to construe the language "denied requests to modify a resource allocation" is not helpful. For reasons detailed in section II.B, IV's alternative construction should be rejected because (i) there is no intrinsic support for this interpretation of this claim language, (ii) a "resource unavailable message" / "resource denial" is clearly different from a "requested modified resource allocation", and (iii) IV's proposal for these terms remains at odds with its own proposed constructions of "modified resource allocation" and "modify a resource allocation."

Under Federal Circuit case law, the Court can only correct a mistake in the claim language if "the correction is not subject to reasonable debate based on consideration of the claim language and the specification." *Novo Indus., L.P. v. Micro Molds Corp.*, 350 F.3d 1348, 1354 (Fed. Cir. 2003). IV has not met that standard here. *Imperium (IP) Holdings, Inc. v. Apple, Inc.*, 920 F. Supp. 2d 747, 753 (E.D. Tex. 2013) (finding claims indefinite because there was no antecedent basis for a plural term "pixels" which was argued to refer back to a singular term "pixel"); *see also Red Rock Analytics, LLC v. Samsung Elecs. Co.*, No. 2:17-cv-101, 2018 WL 1806859, at *18 (E.D. Tex. Apr. 16, 2018). As such, because a POSITA would be unable to discern the scope of this language with reasonable certainty, the Court should hold that this claim term is indefinite.[6]

F. "component configured to" Means-Plus-Function Terms ('686 patent claim 7)[7]

IV's arguments are unpersuasive. First, IV argues that VMware cannot show that "the claim "recites function without reciting sufficient structure for performing that function" because the preamble and element preceding the "component" limitations recite structure such as a "virtual server operating in a first physical host" and "one or more processors and one or more memories." Dkt. No. 53 at 26–27.

As a preliminary matter, IV's argument regarding the effect of the preamble on the remainder of the claim is fundamentally contrary to its stance that the preambles of the claims 1 and 4 of the '726 patent (which also recite a "virtual server operating in a first physical host") are not limiting. Dkt. No. 53 at 15. IV cannot have it both ways.

Notwithstanding this issue, IV fails to acknowledge that the claimed "components" are not recited as part of **any** of the claim terms that it contends provide structure to the claims. Instead,

[6] Although VMware initially proposed an alternative construction of this claim term, the parties' respective opening briefs have made it clear that this alternative cannot resolve the indefiniteness issue present in the '686 patent claims.

[7] The parties proposed terms and construction are identified at Dkt. No. 54 at 21–24.

the claim recites broadly, "[a] system" for modifying a virtual server in a physical host, where the components are merely part of the "system." '686 patent, cl. 7. It is of no moment whether "virtual server," "physical host", "processors" or "memories" connote structure when the claims do not require the claimed component terms to be a part of these alleged structures.

IV's reliance on *Fisher-Rosemount Sys., Inc. v. ABB Ltd.* is misplaced. No. 4:18-cv-00178, 2019 WL 6830806 (S.D. Tex. Dec. 12, 2019). *Fisher-Rosemount* did not involve claim language of "a component configured to" perform a function. Instead, *Fisher-Rosemount* found that the claim term "**processor**" connotes structure because, *inter alia*, a processor is a "tangible object that can be purchased and that can perform certain functions even without specific instructions." *Id*. at *16. Additionally, the court noted that claim "describes how the processor interacts with the invention's other components and identifies where the processor is located." *Id*. In contrast, in the present case, the claim language of "a component configured to" clearly doesn't refer to a tangible object that can be purchased and perform functions without specific instructions. Nor does the claim specify *any* relationship between the claimed components and the rest of the system. As such, "a component configured to" is a nonce term—indeed *Fisher-Rosemount* acknowledged that similar terms (such as "element" and "device") "typically do not connote structure." *Id*.

IV's alternative identification of structure disclosed in the specification also fails. First, IV's citation to Figure 1, 2:63–3:4, 3:59–4:8, 5:7–28, 5:42–62[8] and multiple structures "Dynamic Resource Configuration Module 100; Physical Hosts 160A-C; Virtual Servers 162A-G" for each claimed "component," is unsupported attorney argument. A POSITA reading these sections would not understand them to disclose support for the claimed functions. Dkt. No. 54-2 (Snoeren Decl.)

[8] In the parties' various exchanges, and despite requests from VMware, IV has never cited to portions of the specification that it contends perform the claimed function. IV disclosed these citations for the first time in its opening brief.

at 14–20. In particular, the claimed functions here, for the three claimed component terms, generally relate to (i) receiving an indication that a first host is overloaded based on a specific determination of whether a virtual server is overloaded based on one or more resource unavailable messages resulting from denied requests to modify a resource allocation; (ii) determining that the second physical host can accommodate the requested modified resource allocation; and (iii) generating a host transfer signal indicating a second physical host. '686 patent, cl. 7. IV's specification citations completely fail to disclose a software algorithm for performing the first two claimed functions. For example, the specification fails to discuss *anything* related to the claim language of determining that a virtual server is overloaded based on one or more resource unavailable messages resulting from denied requests to modify a resource allocation.[9] As such, this claim is indefinite.[10]

III. DISPUTED TERMS FROM U.S. PATENT NO. RE42,726 (THE "'726 PATENT")

A. Terms that overlap with disputed claim terms in the '686 patent

Similar / Overlapping Terms		
Term	**VMware Proposal**	**IV Proposal**
"modify[ing] a resource allocation for the virtual server" / "modifying [the] computer resources allocated to a virtual server" ('726 patent claims 1, 4-5 and 8)	"modify[ing] a quality of service guarantee for the virtual server" / "modifying [the] quality of service guarantee of a virtual server" *See also* construction of "quality of service guarantee"	"modify[ing] a set of the functions and features of a physical host(s) used in implementing tasks for each virtual server" / "modify[ing] a set of the functions and features of a physical host(s) used in implementing tasks for each virtual server"

[9] Additionally, IV's citations fail to discuss anything related to the claim language of determining that a virtual server is overloaded based on one or more resource unavailable messages resulting from denied requests to modify a resource allocation.

[10] As to the third claimed function, IV's specification citations miss the mark because they fail to include citations to the algorithm in the specification describing "Dynamic Virtual Server Mover 140." While column and lines 5:7–28 (cited by IV) discuss the virtual server mover at a high level, the algorithm corresponding to the Figure 6 (and described in the specification at 12:1–28) is necessary to provide structure to this means-plus-function term.

"virtual server overloaded signal" ('726 patent claims 1, 4-5 and 8)	"signal indicating that an average number of resource denials for a virtual server is beyond a pre-configured threshold" *See also* construction of "resource denials"	"an indication that a virtual server has been or is being denied resources"
"virtual server" ('726 patent claims 1-11)	"a process executing on a host computer that accepts communication requests, and that is capable of receiving a quality of service guarantee from a physical host"	plain and ordinary meaning; in the alternative: "a virtual machine that resides on a physical server and uses the physical server's resources but has the appearance of being a separate dedicated machine"

For the reasons set forth in sections II.A, II.C, and II.D that respond to IV's arguments regarding these terms, the Court should adopt VMware's proposed constructions of these terms.

B. "resource denials" ('726 patent claims 1, 4–5 and 8)

VMware Proposal	IV Proposal
"indications that requests by the virtual server for additional resources are either implicitly or explicitly denied"	"an indication that a request by the virtual server cannot be immediately serviced"

IV contends that its proposed construction "is a direct quotation from the portion of the specification where the invention as a whole […] is described" while "VMware's proposed construction is merely one narrow example of a single preferred embodiment." Dkt. No. 53 at 15. Not so.

First, as detailed in VMware's opening brief, the parties have reached substantial agreement regarding the first part of the construction of "resource denials" – i.e., "indications that requests by the virtual server." Dkt. No. 54 at 24–25. But, a subtle, but important, distinction for the present term is that VMware incorporates plural nouns in the agreed language as "indications…" to account for the fact that the claim term resource **denials** is a plural noun. In

contrast, IV improperly seeks to alter this term to singular tense as "an indication." Dkt. No. 54 at 24–26.

Second, unlike VMware's proposed construction, IV's proposal omits the word "resource" from the claim term, improperly attempting to broaden out the requests by the virtual server for additional resources to simply any request by the virtual server. These attempts to substantially re-write the claims are contrary to black letter claim construction law and should be rejected.

Third, VMware's proposal is taken from a definitive statement in the specification of "resource denials," which explains that resource denials "**are**" indications that requests by the virtual server "**for additional resources [are] either implicitly or explicitly denied**." '726 patent, 7:51–54 (emphasis added). By contrast, IV's proposal has been selectively taken from language that discusses what a resource denial "**may refer to**." '726 patent, 2:55–61 (emphasis added). Moreover, IV relies upon a section of the specification that refers to a specific embodiment. *See* '726 patent at 2:52–61 ("In one embodiment … [a] resource denial may refer to any request by the virtual server that cannot be immediately serviced").

Fourth, IV's argument that VMware's proposed construction is based on an illustrative embodiment is not correct. Neither the sentence nor the paragraph in the specification from which VMware's proposed construction is drawn uses the language "embodiment." *See* '726 patent, 7:50–64. And the sentence from which VMware's proposal is drawn makes it clear that the characterization of resource denials was intended to be definitional. By contrast, as noted above, the phrase "[i]n one embodiment" is used in the same paragraph, and indeed, in the sentence preceding the one from which IV draws its proposed construction. Under IV's argument, this fact undermines its own proposed construction. For these reasons, IV's proposed construction should be rejected and VMware's proposed construction should be adopted.

C. "quality of service guarantee" ('726 patent claims 1 and 4)

VMware Proposal	IV Proposal
"information that specifies a guaranteed amount of an assigned resource, and that can be dynamically increased/modified"	"a guaranteed resource allotment which can be dynamically increased/modified"

IV takes issue with two portions of VMware's proposed construction – "information that specifies" and "an assigned resource" – and contends that "the intrinsic record is devoid of any such limitation." Dkt. No. 53 at 20.

As to the language "information that specifies," IV contends that the claim term "quality of service guarantee" is not information and is different from "quality of service guarantee information." This argument obfuscates what is a simple issue. A quality of service guarantee is unquestionably information. The specification is replete with examples that make it clear that this must be the case.

For example, the summary of the invention states that "[t]he present invention **dynamically adjusts the quality of service guarantees** for virtual servers based upon the resource demands experienced by the virtual servers." '726 patent, 2:38–40 (emphasis added). It is unclear, and IV has not explained, how it would be possible to dynamically adjust quality of service guarantees if they did not refer to information. The specification also states that "[t]he term "virtual server" as used herein refers to a virtual server capable of **receiving a quality of service guarantee** from a physical host." '726 patent, 3:43–45 (emphasis added). It is unclear how a virtual server would receive a quality of service guarantee from a physical host if it was not information. The specification goes on to state that "[a] **resource allocation for a virtual server is specified as a "quality of service guarantee"** for that particular server. Each **physical host stores quality of service guarantees** for the virtual servers it hosts." '726 patent, 4:39–42 (emphasis added). Again, it is not clear how a resource allocation could be specified as a quality of service guarantee or how

a physical host could store quality of service guarantees if these terms did not refer to information. These instances are provided by way of example and without limitation, as there are other such examples in the specification. As such, IV's arguments as to the language "information that specifies" should be rejected and VMware's proposed construction, which properly incorporates this language into the construction, should be adopted.

As to the language "assigned resource," IV argues that "nowhere in the specification is a quality of service guarantee described as being limited to 'an assigned' resource." Dkt. No. 53 at 21. This is not true. The specification states:

> A customer's virtual server is typically **assigned a fixed level of resources**, **corresponding to either a fixed percentage of the capacity of a particular physical host** (for example, the operating system may be instructed to allocate twenty percent of the central processing unit cycles to process A and two percent to process B) **or a fixed number of units** (for example, the operating system may be instructed to allocate X cycles per second to process A and Y cycles per second to process B).

'726 patent, 2:5–13 (emphasis added). This directly parallels later language that states:

> In one embodiment, each individual virtual server 162 has a different quality of service guarantee. Different **quality of service guarantees are implemented by allocating different amounts of the resources of each physical host machine** 160 to servicing each of the virtual servers 162. **Physical host 160 resources may be allocated as percentages of the resources of a particular physical host 160, or as a particular number of units within a physical host 160** (for example, the operating system may be instructed to allocate X cycles per second to process A and Y cycles per second to process B)

'726 patent, 3:66–4:8 (emphasis added). As such VMware's proposed language finds clear support in the specification. By contrast, the language "guaranteed resource allotment," and sub-phrases "resource allotment" and "allotment" (and for that matter, other versions of allotment such as "allot" or "allotted"), in IV's proposed construction are not found *anywhere* in the specification. As such, for the reasons IV's itself sets forth, its proposed construction should be rejected in favor of VMware's proposed construction.

D. Mean-Plus-Function Elements ('726 Patent claims 1, 3, 4, 5, 7)

VMware contends that each of the terms identified in Dkt. No. 54 at 30–33 should be construed under § 112 as means-plus-function terms. IV in contrast argues that terms recited in the preamble such as "network system," "physical hosts," "virtual server," and "computer resources," "give context to the environment in which the 'virtual server resource monitor' is claimed." Dkt. No. 53 at 28. As discussed in section II.F, this is directly contradicted by IV's argument that the preambles of the '726 patent claims are not limiting. Furthermore, as discussed in section II.F, none of these terms are recited as part of **any** of the claim terms that IV contends provide structure. Instead, the '726 patent claims 1, 3, 4, 5, 7 recite a "network system," "computer program," or "system" comprising these separate claim terms and do not provide context or structure **any** of these terms. VMware's opening brief provides additional discussion regarding why these are black-box terms that fail to carry structural meaning to a POSITA. *See e.g.,* Dkt. No. 54 at 27–33; Dkt. No. 54-2 (Snoeren Decl.) at 24–42. Additional reasons why the Court should not adopt IV's proposals for these terms are discussed below.

a. "a virtual server resource monitor [communicatively coupled to the first physical host and] configured to monitor resource denials and to send a virtual server overloaded signal in response to the resource denials" ('726 patent claims 1 and 5) // "program code for creating a virtual server resource monitor communicatively coupled to the first physical host and configured to monitor resource denials and, in response to the resource denials, to send a virtual server overloaded signal" ('726 patent claim 4)

Every single instance of this term in the specification describes what the black-box "monitor" *does*, not what it *is*. Because the term itself does not carry structural meaning to a POSITA (either based on or independent of the intrinsic record), and because the claims functionally recite what this black-box module is "configured to" do, this term invokes § 112 ¶ 6. Also, as described in VMware's opening brief, the term "program code" does not on its own

connote structural meaning to a POSITA, and is not discussed anywhere in the specification.[11] As such, and because the claims functionally recite what this black-box term is "configured to" do, this term invokes § 112 ¶ 6.

IV's alternative identification of structure disclosed in the specification also fails. IV cites to Figure 1, Figure 3, 4:64–5:4, 5:44–65[12] and multiple structures including "dynamic resource configuration module 100 includes, among other things, virtual service resource modifier, which in combination with physical hosts 160A-C monitors resource denials and sends a virtual server overload signal" as providing structure for the "virtual server resource monitor." However, IV specifically does not cite to the "virtual server resource monitor," even though it separately contends that this term is sufficiently described in the specification to connote structure for a POSITA. Dkt. No. 53 at 29–30. This argument does not make sense.

By contrast, VMware points to the software algorithm in the specification that describes the functions of the "Virtual Server Resource Monitor 110," which is the same as the claim term, as the disclosure of structure. VMware also points to a more complete discussion of this term in the specification than IV, which specifically omits column and lines 7:41–9:46 despite its disclosure of a software algorithm that performs the claimed function. IV never explains how or why VMware's identification of structure unreasonably limits the claimed functionality. If anything, IV's limited citations that fail to include the relevant algorithm, are incomplete.

b. "a virtual server resource modifier [communicatively coupled to the first physical host and] configured to receive the virtual server overloaded signal

[11] IV's reliance on the *RLIS*, *Eolas*, and *Aloft* opinions does not warrant a different result. Dkt. No. 53 at 29. Each of these cases were decided under the overruled "strong presumption" standard. *Compare RLIS, Inc. v. Allscripts Healthcare Solutions, Inc.*, Nos. 3:12cv208, -209, 2013 WL 3772472, at *14 (S.D. Tex. July 16, 2013) ("[T]he presumption flowing from the absence of the term 'means' is a strong one that is not readily overcome") with *Williamson*, 792 F.3d at 1349 ("expressly overrule the characterization of that presumption as 'strong'").

[12] Again, IV disclosed these citations for the first time in its opening brief.

and, in response to the virtual server overloaded signal, to modify a resource allocation for the virtual server and to send a virtual server resource modification signal" ('726 patent claims 1 & 5); "program code for creating a virtual server resource modifier communicatively coupled to the first physical host and configured to receive the virtual server overloaded signal and, in response to the virtual server overloaded signal, to modify a resource allocation for the virtual server and to send a virtual server resource modification signal" ('726 patent claim 4)

For the reasons discussed above in section III.D.a, the Court should construe the claim language of "a virtual server resource modifier … configured to … [perform a function]" and "program code for [performing a function]" as invoking 35 U.S.C. § 112 ¶ 6. Every single instance of this term in the specification describes what the black-box "modifier" *does*, not what it *is*.

IV's alternative identification of structure disclosed in the specification also fails. IV cites to Figure 1, Figure 2A, Figure 4, 4:64–5:20[13] and multiple structures including "dynamic resource configuration module 100 includes, among other things, virtual service resource modifier, and in combination with physical hosts 160A-C and virtual servers 162A-G receives virtual server overload signals and signals a resource modification is needed" as providing structure for the "virtual server resource modifier." This argument and the identification of other multiple other components than "Virtual Server Resource Modifier 120" for performing the functions of the claimed "virtual server resource modifier" does not make sense.

By contrast, VMware points to the software algorithm in the specification that describes the functions of the "Virtual Server Resource Modifier 120," which is the same as the claim term, as the disclosure of structure. VMware also points to a more complete discussion of this term in the specification than IV, which specifically omits column and lines 9:47–10:52 despite its disclosure of a software algorithm that performs the claimed function. IV never explains how or

[13] Again, IV disclosed these citations for the first time in its opening brief.

why VMware's identification of structure unreasonably limits the claimed functionality. IV's limited citations that fail to include the relevant algorithm, are incomplete.

c. "a load balanc[ing/er] [module] [communicatively coupled to the plurality of physical hosts and] configured to receive the virtual server resource modification signal and to determine whether the first physical host is overloaded and, in response to a determination that the first physical host is overloaded, to send a physical host transfer signal that indicates a second physical host" ('726 patent claims 1 and 5) // "program code for creating a load balancing module communicatively coupled to the plurality of physical hosts and configured to receive the virtual server resource modification signal and to determine whether the first physical host is overloaded and, in response to a determination that the first physical host is overloaded, to send a physical host transfer signal that indicates a second physical host" ('726 claim 4)

IV contends that the terms "load balance[r] module" and "program code for creating" the same are "specific and well-known construct[s] with a structure known by those of skill in the art." Dkt. No. 53 at 32. IV provides no support for this argument. A POSITA would not understand these terms to refer to well-known structures. Dkt. No. 54-2 (Snoeren Decl.) at 33–34. Indeed, IV even admits that these terms relate broadly to something "typically implemented in software." Dkt. No. 53 at 32. IV's descriptions of this module show that it invokes § 112(6). *Id*. at 32–33. *See, e.g., Sound View Innovations, LLC v. Facebook, Inc*., No. 16-cv-116 (RGA), 2017 WL 2221177, at *2 (D. Del., 2017) ("For software patents claiming a function that a general purpose computer cannot perform, the specification must disclose an algorithm."). Indeed, similar to the previous terms, every single instance of this term in the specification describes what the black-box "module" *does*, not what it *is*. Because the term itself does not carry structural meaning to a POSITA (either based on or independent of the intrinsic record), and because the claims functionally recite what this black-box module is "configured to" do, this term invokes § 112 ¶ 6.

IV's alternative identification of structure disclosed in the specification also fails. IV cites to Figure 1, Figure 5, Figure 6, 4:64–5:20, 10:53–11:20[14] and multiple structures including "dynamic resource configuration module 100 includes, among other things, load balancer, and in combination with physical hosts 160A-C and virtual servers 162A-G receives virtual server resource modification signal and determines whether the physical host is overloaded and sends a host transfer signal if it is" as providing structure for the "load balancing module." This argument and the identification of other multiple other components than "Physical Host Load Balancing Module 130" for performing the functions of the claimed "load balancing module" does not make sense.

By contrast, VMware points to the software algorithm in the specification that describes the functions of the "Physical Host Load Balancing Module 130," which is effectively the same as the claim term, as the disclosure of structure. VMware also points to a more complete discussion of this term in the specification than IV, which specifically omits, among other cites, column and lines 11:21–52 despite its disclosure of a software algorithm that performs the claimed function. IV never explains how or why VMware's identification of structure unreasonably limits the claimed functionality. IV's limited citations that fail to include the relevant algorithm, are incomplete.

d. "a dynamic virtual server mover [communicatively coupled to the plurality of physical hosts and] configured to receive the physical host transfer signal and, in response to the physical host transfer signal, to transfer the virtual server from the first physical host to the second physical host" ('726 patent claims 1 and 5) // "program code for creating a dynamic virtual server mover communicatively coupled to the plurality of physical hosts and configured to receive the physical host transfer signal and, in response to the physical host transfer signal, to transfer the virtual server from the first physical host to the second physical host" ('726 patent claim 4)

[14] Again, IV disclosed these citations for the first time in its opening brief.

For the reasons discussed above in section III.D.a, the Court should construe the claim language of "a dynamic virtual server mover … configured to … [perform a function]" and "program code for [performing a function]" as invoking 35 U.S.C. § 112 ¶ 6. Every single instance of this term in the specification describes what the black-box "mover" *does*, not what it *is*.

IV's alternative identification of structure disclosed in the specification also fails. IV cites to Figure 1, Figure 6, 4:64–5:20, 6:20–31, 6:39–59, 12:1–10[15] and multiple structures including "dynamic resource configuration module 100 includes, among other things, dynamic virtual service mover and in combination with physical hosts 160A-C" as providing structure for the "dynamic virtual server mover." This argument and the identification of other multiple other components than "Dynamic Virtual Server Mover 140" for performing the functions of the claimed "dynamic virtual server mover" does not make sense.

By contrast, VMware points to the software algorithm in the specification that describes the functions of the "Dynamic Virtual Server Mover 140," which is the same as the claim term, as the disclosure of structure. VMware also points to a more complete discussion of this term in the specification than IV, which specifically omits portions of column and lines 11:63–12:23 despite its disclosure of a software algorithm that performs the claimed function. IV never explains how or why VMware's identification of structure unreasonably limits the claimed functionality. IV's limited specification citations that fail to include the relevant algorithm, are incomplete.

e. "the dynamic virtual server mover is further configured to direct the first physical host to store, in the file system, a set of system files for the virtual server and to direct the second physical host to access, from the file system, the set of system files for the virtual server, thereby transferring the virtual server from the first physical host to the second physical host" ('726 claims 3 and 7)

[15] Again, IV disclosed these citations for the first time in its opening brief.

For the reasons discussed above in section III.D.d, and in VMware's opening brief, the Court should construe the language which recites "dynamic virtual server mover … configured to … [perform a function]" as invoking 35 U.S.C. § 112 ¶ 6.

IV. DISPUTED TERMS FROM U.S. PATENT NO. 7,949,752 (THE "'752 PATENT")

A. "exhausted" ('752 patent claims 1, 9 and 24)

VMware Proposal	IV Proposal
"unavailable for reuse"	"used up to the allotted or pre-determined amount"

IV makes two primary arguments for this term. Neither are persuasive. And tellingly, IV has failed to address the clear prosecution disclaimer supporting VMware's proposed construction. *See, e.g.,* Dkt. No. 53 at 9–10.[16]

First, IV argues that the intrinsic record is "replete with evidence supporting IV's proposed construction." Dkt. No. 53 at 9 (citing various passages from the '752 patent's specification). Not so. The '752 patent only uses the word "exhausted" once—specifically in the context of halting consumption of a service resource when "the amount held by an agent is exhausted." '752 patent at 25:30–34. None of the other sections cited by IV reference the concept of "exhausted." Moreover, there are no passages in the specification that overcome IV's clear and unambiguous prosecution disclaimer labeling the "exhausted" claim term as the **exact opposite** of **reusable** resources. Dkt. No. 54-14 (2010-08-20 Response) at 14 (emphasis added).

[16] Furthermore, IV's argument that the Court should follow Magistrate Judge Mitchell's prior report and recommendation ("R&R") for this term to promote "uniform treatment of claim construction" is inapposite. Dkt. No. 53 at 9. That R&R was merely a "recommendation" and was not adopted by the District Court Judge before the case settled. Indeed, **both parties** objected to the Magistrate's recommendations before settling. Because the claims were not ultimately construed by the District Judge, there exists no risk of inconsistent rulings on the terms.

The bulk of IV's remaining argument is that IV's proposal reflects the overall goal of the invention. Again, this is not the case. VMware's proposal is consistent with the invention disclosure, while IV's proposal incorrectly interprets a particular embodiment and then seeks to read this interpretation into the claims. *See, e.g.*, *GE Lighting Sols., LLC v. AgiLight, Inc.*, 750 F.3d 1304, 1309 (Fed. Cir. 2014).

The only paragraph that discusses the concept of exhaustion ('752 patent, 25:20–34) illustrates why VMware's proposal is correct. This paragraph is written with reference to "method 800" which, in relevant part, includes steps 810 and 820. *Id.* at 25:20; *see also id.* at Fig. 17 (depicting method 800). Step 810 performs an evaluation of an assigned "service permission" in order to "identif[y] the amount of a service resource 25 which is allocated to agent 22." *Id.* at 24:56–58; *see also id.* at 18:19–21 ("a service permission can specify a pre-authorized amount of service resource 25 which is allowably consumed by the respective agent"). After executing the service, step 820 then "decrement[s] the amount allocated to agent 22 by the amount actually used." *Id.* at 25:16–19; *see also* Fig. 17 ("decrement A_h by A_u"). VMware's proposal is consistent with this disclosure because it recognizes that once the service's resources (A_h) have been fully decremented/used up, they are not reusable without issuance of an additional service permission assigning a new A_h amount of service resources.

IV's proposal, in contrast, is inconsistent with this description. IV's proposal contemplates that the service resource should be available again after the customer "has released the resource." Dkt. No. 53 at 5. But method 800 does not include a step where the amount of service resource A_h is **incremented** by an amount of service resource released. Nor is that process discussed anywhere else in the '752 patent. Nor is IV correct in arguing that under VMware's proposal "[o]ne customer using an amount of a service and service resource would permanently make that

service and service resource unavailable to subsequent customers." *Id*. The claims don't recite that the entirety of a service, and/or a service resource, applicable to multiple customers is exhausted. Instead, they recite a different customer-specific concept, of "a service resource **configured to be consumed by the network-based agent**" where "**an amount** of the service resource is exhausted." *See, e.g.,* '752 patent cl. 7. As discussed above, and in VMware's opening brief, the specification also repeatedly and consistently describes this process as customer specific.

IV's proposal is also inconsistent with the two examples of service resources discussed with connection to method 800: "an amount of storage space" (for use with a voicemail service provider) and an amount of "connect time" (for a long distance service provider). '752 patent at 25:20–34. For example, claim 1 becomes incomprehensible when using IV's proposals for the "exhausted" and "consumption" terms, reading as: "wherein an amount of the service resource is ~~exhausted~~ used up to the allotted or pre-determined amount upon being ~~consumed~~ used by the network-based agent." As written, the user's apportioned resource would be "used up" immediately upon being used. This does not make sense, and it is directly contrary to IV's own articulation of the overall goal of the invention: allowing a customer to pay for computing power, time and access on an as needed basis. Dkt. No. 53 at 9–10. VMware's proposals, in contrast, render the claim viable, reading as: "wherein an amount of the service resource is ~~exhausted~~ unavailable for reuse upon being ~~consumed~~ used up by the network-based agent."

B. "consumed" (recited in '752 patent claims 1, 9 and 24)

VMware Proposal	IV Proposal
"used up"	"used"

As articulated in its opening brief, VMware's proposal for this term is supported by narrowing statements made during prosecution and a definitional statement provided in the patent specification. As evidenced by these statements, "consumed" cannot simply mean "use" as IV

proposes—indeed, the claim language was amended to employ both "using" and "consumed" evidencing that the two words have different meaning. *See, e.g.*, Dkt. No. 54-12 (2009-11-13 Amendment) at 5–6. The applicant then used this narrowing amendment to argue around prior art that disclosed an agent configured to use a resource. *Id.*, at 11 ("Humpleman does not teach or suggest 'a service and a service resource **configured to be consumed by the agent** . . . wherein the service resource is **exhausted** after it is consumed by the agent," as recited by claims [sic] 86, 92, 94, 106, and 107.") (emphasis retained).

Indeed, IV fails to substantively address the narrowing and definitional statements in the intrinsic record (also identified in VMware's opening brief). Instead, IV only points to three areas of the '752 patent, none of which provide a definitional statement of the term "consumed" or otherwise disclose anything other than "consumed" and "used" being included in sentences using the terms as alternative options. *See,* Dkt. No. 53 at 10.

IV cannot escape the fact that the terms "exhausted" and "consumed" are intertwined. VMware's proposed constructions are consistent with the specification and file history and also are consistent with one another. On the other hand, IV's constructions for "exhausted" and "consumed" when read together make the claim language incomprehensible.

C. "service" ('752 patent claims 1, 3, 9 and 24)

VMware Proposal	IV Proposal
"An application that is used by an agent on behalf of a principal"	"Network functionality available to agent(s)/network-based agent(s)"

VMware's proposal is consistent with the specification's disclosure that an "agent uses the service on behalf of a principal" and that the services "comprise one or more software applications providing various capabilities that are available to a principal." '752 Patent at 3:3–5; 10:17–19;

see also id. at 10:19–21 ("Each service 24 may be utilized by one or more agents 22 in order to perform their respective tasks.").

IV's proposal, in contrast, does not help clarify the claim scope. IV cites to the abstract's disclosure that "an agent is operable to utilize a service within the network system"—but IV appears to only cite this section as support that an agent used a service in a **network system**. This clarification, however, is not necessary because the claims already recite that the agent uses a service in a **network system**. *See, e.g.*, '752 patent claim 1 ("means, including the network-based agent, for using a service"). Also, IV's proposed language of "network functionality" could be interpreted as relating to functionality of a network, such as signal processing and transmission—a concept that is inconsistent with the '752 patent's more general description of service applications that utilize a network system but are not network functionality itself, such as "an e-mail service, a voice mail service, a paging/facsimile service, an address book and calendar service, and a business news and stocks information service." '752 patent at 13:25-28.

Nor is IV correct in arguing that the patent discloses examples of the claimed service as "sub-systems" and not "applications" (Dkt. No. 53 at 11). Instead, the specification actually describes "sub-systems" (e.g., "mass storage subsystem[s] of tapes or disk drive" *see* '752 patent at 13:49) as something that ***supports*** a service, but not as the service itself. *See id.* at 13:20–26 ("At least a portion of the sub-systems in computer-based system 30 may support one or more services 24").

Finally, IV's own dictionary definition contradicts its proposal. This definition states that "services" are "specialized, **software-based** functionality provided by network servers." Dkt. No. 53 at 12 (emphasis added). While IV emphasized that the definition includes the word

"functionality," the preceding "specialized, software-based" portion of the definition (i.e., an application) is glossed over.

D. Means-Plus-Function Terms

VMware's proposals for these "means" terms identify "the corresponding structure … described in the specification" as required by 35 U.S.C. § 112. For example, for the "means for mediating an interaction between the means for using the service and the service" term in claim 3, VMware proposes that the structure is a "service wrapper (26) as described in [the '752 patent at] 16:22–38." The specification citation refers to a passage titled "Service Wrapper (Details))" which provides a description of the structure for "service wrapper (26)." This section describes that "service wrapper (26)" is composed of "a converter 48 and a monitor 50." *Id*. at 16:28–29.

In contrast, IV's proposal ("service wrapper (26)"), should be rejected because it does not identify the corresponding structure described in the specification. IV apparently intends to leave it up to the jury to determine whether the disclosed structure for this means term should be interpreted to require, e.g., the "converter 48" and "monitor 50a" as disclosed in the specification. This is improper—construing claim scope is the province of judges. *Markman v. Westview Instruments, Inc*., 517 U.S. 370, 372 (1996). Moreover, imposition of IV's vague proposal would render the claim indefinite. *Compare in re TLI Comm'cns LLC Patent Litig.*, 87 F.Supp.3d 773, 800 (E.D. Va., 2015) ("simply disclosing a black box that performs the recited function is not a sufficient explanation of the algorithm required to render the means-plus-function term definite") (quotations omitted).

VMware did not "cherry-pick[]" portions of the specification ignoring other structure for this term as IV argues. Dkt. No. 53 at 7. Most of IV's additional citations merely discuss a "service wrapper (26)" in functional terms devoid of structure. *See* '752 patent at 3:20–27 (function: "cooperat[ing] with the agent server to mediate interaction"); 17:43–46 (function: "control[ing]

access"); 18:49–54 (no statement of structure or function). The only possible exception is IV's citation to the '752 patent at 25:1–24. Dkt. No. 53 at 13. This section describes an algorithm for service wrapper 26 including steps 802-820. To the extent this Court finds this disclosure as structure, VMware offers a revised proposed construction of "service wrapper (26) as described in 16:22–38 and 25:1–24." *See Micro Chem., Inc. v. Great Plains Chem. Co.*, 194 F.3d 1250, 1258 (Fed. Cir. 1999) ("When multiple embodiments in the specification correspond to the claimed function, proper application of § 112, ¶ 6 generally reads the claim element to embrace each of those embodiment.").

Finally, the parties dispute the structure for the term "means for monitoring an amount of the service resource used by the network-based agent" of claim 4. VMware's proposal is that the "means for monitoring" is a "monitor as described in '752 patent at 16:50–61" and should be adopted. IV, in contrast, argues that the structure should be taken up to the software package of the "service wrapper 26" which, as described above, is described as including a monitor and a converter. However, IV fails to explain how the converter is at all involved in the claimed function and should be rejected for this reason alone. The converter does not perform the function of monitoring, instead, it is used to "convert between a computer language (or instruction set) used within agent server 20 and a computer language (or instruction set) used within the respective service 24." *See, id*. at 16:30–33.

V. DISPUTED TERMS FROM U.S. PATENT NO. RE43,051 (THE "'051 PATENT")

A. "virtual server" ('051 patent claims 1, 3 and 6)

VMware Proposal	IV Proposal
"a process executing on a host computer that accepts communications requests"	Plain and ordinary meaning, or alternatively: "virtual machine(s) that reside(s) on a physical server and use(s) the physical server's resources but has/have the appearance of being a separate dedicated machine(s)"

IV is right that the term "virtual server" is used in the '051 patent (and in the '686 and '726 patents) consistent with its customary usage at the time of the invention. But IV is wrong that the customary use at the time would have meant a virtual machine (which IV defines as a functional simulation of a computer and its associated devices) instead of the application servers actually described in the intrinsic record. There is no intrinsic support for IV's contention that virtual server means simulated machines, like computers. Instead, consistent with VMware's position, the intrinsic record describes the virtual server of the invention as process based.

For instance, the '051 patent explains that instead of dedicating individual physical host computers to individual customers, a "service provider may utilize one physical host computer to provide commercial host services to multiple customers" using "***a server application*** executing on a single physical host [that] can be programmed to process requests made to multiple network addresses." '051 patent at 2:47–53 (emphasis added). The specification describes this to mean "each customer is assigned a network address (or domain name), and is provided with resources on a single, physical host computer, effectively sharing the host with other customers." *Id.*

The '051 patent further explains how to create these private virtual servers by pointing to an incorporated reference that issued as U.S. Patent No. 6,976,258. *Id.* at 4:64–67. The '258 patent in turn describes a server as "a process, executing on a dedicated physical services client [that] services client requests for a single network address (physical host) only," Ex. 48 at 3:23–25, and explains that these processes can provide for virtual hosting (*i.e.*, servicing requests for multiple network addresses on a single physical host) by "creat[ing] child processes to service the requests," *id.* at 1:24–44. This ability for a single application to create child processes to handle client requests on a plurality of network addresses is what provides for the *virtual* aspect of the

virtual servers of the invention, so that the application can function as multiple servers. *See id.* at 1:45–2:5; Dkt. No. 54-2 (Snoeren Decl.), 43–44.

In contrast to the many discussions of process-based virtual servers in the intrinsic record, there is no support whatsoever for IV's incorrect assertion that virtual server means virtualized computer hardware.[17] IV mistakenly points to disclosures about the virtual servers using resources on the host machine as evidence that these virtual servers are somehow partitioning resources for use in virtualized hardware. The specification, however, is clear that the physical host's resources are being allocated to server *applications*, not virtualized hardware. It explains, for example, that "virtual hosting" involves "a server application executing on a single physical host," '051 patent at 2:47–53, such that each customer "is provided with resources on a single, physical host computer, effectively sharing the host with other customers," *id.* at 2:54–64. In other words, IV's citations to discussions about sharing physical resources on a single host apply to process-based virtual hosting, and do not indicate that the virtual servers of the invention are virtualized hardware, as IV suggests.

IV's citation to two cherry picked dictionary definitions does not trump the intrinsic record's support of VMware's proposal. Moreover, a more thorough review of the extrinsic record reinforces the conclusion that the invention describes process-based virtual servers, not simulated hardware. For instance, Ben and Peter Laurie's book, *Apache: The Definitive Guide* (1999), describes how to configure Apache virtual servers and discusses virtual servers in terms of processes. Dkt. 54-29 at 7, 10, 13–14, 16; Dkt. No. 19 at 6 (emphasis added) ("Apache is an example of a preforking server," which "means that the main server ***starts a pool of processes*** to

[17] *See, e.g.,* Dkt. No. 53 at 47 (IV arguing: "a virtual server, i.e., a software abstraction of physical hardware").

handle client requests, rather than forking a new process for each incoming request."). Moreover, the process-based Apache virtual server is discussed in terms of preforking—which is exactly how the intrinsic record for IV's patent discusses virtual servers of the invention. Ex. 48 ('258 patent) at 1:45–2:5; *see also* Dkt. No. 54-2 (Snoeren Decl.) at 42–45.

Given the background provided in the specification, including about the function of forking child processes to service client requests on multiple network addresses, a process-based server application like Apache is what a POSITA would have understood the term virtual server to mean at the time of the invention. As a result, the Court should reject IV's arguments that attempt to read virtualized hardware, like virtual machines, into the claims, and instead give the term virtual server its plain meaning as evidenced by the intrinsic and extrinsic record: a process executing on a host computer that accepts communications requests.

B. "physical interface[s]" ('051 patent claims 1 and 3)

VMware Proposal	IV Proposal
"hardware that provides a point of communication between two or more devices"	Plain and ordinary meaning.

IV argues that physical interface should be given its plain and ordinary meaning, but IV avoids the actual issue: the parties have a fundamental dispute about what the plain and ordinary meaning of physical interface is. IV has steadfastly refused to confirm whether it agrees (or disagrees) that the term "physical" in "physical interface" refers to hardware. IV's opening brief continues to avoid taking a position on whether physical means hardware, and instead focuses on whether the described interface provides a point of communication between two or more devices.

This argument is a red herring because it ignores the crux of the parties' dispute. Moreover, IV is wrong that a physical interface is not a point of communication between two or more devices. Independent claims 1 and 3 both describe the method as "providing private network services . . . in

a location remote from private network users." Reading claims 1 and 3, there is no doubt that the physical interface is providing a point of communication between at least the remote location and the private network users. IV incorrectly tries to sidestep this by pointing out that the tunnel switch may have only one physical interface. But even if true, that does not change that the claims require the physical interface receive transmissions from, and send transmissions to, other devices. *See, e.g.*, '051 patent at claim 1 ("receiving . . . a transmission on a physical interface" and "sending the transmission . . . on the determined physical interface"). There is no indication, even in IV's proposed embodiment where the tunnel switch has only one physical interface, that the physical interface would receive transmissions from itself and send transmissions to itself, particularly given the explanation in the preamble that the transmissions relate to providing network services between network users and a remote location. In other words, IV's argument on this claim term is at best a distraction from the actual dispute between the parties.

Again, the real dispute between the parties is whether the physical interface is hardware, and IV is wrong that the Court should simply disregard this dispute under the guide of applying plain and ordinary meaning. The Court should rather give real meaning to the patentee's choice to use "physical interface" instead of just "interface". Considering how the word physical is used in the specification, how a POSITA would understand that term in the claims is clear: it differentiates the interface as hardware, not a virtual or logical device. For example, the specification uses physical to differentiate between the "physical host" and the "virtual servers" running on it. *See, e.g.*, *id.* at 2:47–64, 4:35–38, 4:56–65, 7:1–11, 7:47–48, 8:4–5. It also explains that the tunnel is a logical construct by differentiating it from a "physical connection." *See id.* at 10:17–20. Although IV sidesteps this entirely, given this context, a POSITA would have understood that the term physical interface refers to hardware; and as explained above, more specifically hardware that

provides a point of communication between two or more devices. The Court should therefore adopt VMware's proposal and reject IV's call for the Court to avoid resolving this fundamental dispute between the parties.

C. physical interfaces and tunnel identifiers in the storing / receiving / determining / sending terms ('051 patent claims 1 and 3)

VMware's proposal adds clarity to the relevant claims, which otherwise risk confusing the jury through recitals of multiple physical interfaces and tunnel identifiers at various steps. *See, e.g.*, '051 patent at cl. 1. For example, some physical interfaces and tunnel identifiers are stored in a customer lookup table. *Id.* Other physical interfaces and tunnel identifiers are stored in customer forwarding tables. *Id.* A transmission is received on a physical interface containing a tunnel identifier, and later, a transmission is sent on a physical interface with a tunnel identifier. *Id.* The claims as written are vague as to which tunnel identifiers and which physical interfaces are being used at each step. But the specification provides clarification: the physical interfaces and tunnel identifiers associated with the receiving step (and thus the customer lookup tables) are the *incoming* physical interfaces and the *incoming* tunnel identifiers, whereas the physical interfaces and tunnel identifiers associate with the sending step (and thus the customer forwarding tables) are the *outgoing* physical interfaces and the *outgoing* tunnel identifiers. *See id.* at 11:30–32, 11:64–12:1, 12:10–20, 12:30–33, 12:59–63, 13:5–11 Figs. 8–9.

IV is wrong to argue that VMware "attempt[s] to limit the disputed terms to a single directionality requirement." Dkt. No. 53 at 51. VMware's proposal does not speak to whether a tunnel is uni- or bi-directional, and instead merely clarify which of the "physical interfaces" and "tunnel identifiers" are being recited. *See* Dkt. No. 54-8. These proposals are necessary because the claims vaguely recite more than one "***a*** physical interface" and tunnel identifier. *See, e.g.*, '051 patent at cl. 1.

IV is also wrong to argue that VMware's proposal contradicts the specification's teaching of a tunnel switch with "one or more physical interfaces." *See* Dkt. No. 53 at 51 (emphasis in original). There is no implication in VMware's proposal that the *incoming* physical interface and the *outgoing* physical interface cannot be the same. This would be intuitive to a POSITA, who would understand that typical network cards are bidirectional and can function both as an outgoing interface and as an incoming interface. IV is also wrong to criticize VMware's proposal for using the word "pair." *See* Dkt. No. 53 at 51. VMware's proposals do not use the word "pair". *See* Dkt. No. 54-8.

VMware's proposal clarifies these terms in a manner consistent with the claim language and specification. The *incoming* tunnel identifiers and the *incoming* physical interfaces are associated with the *receiving* step in the claims. There is no reasonable reading of the claims where the receiving step is not dealing with an incoming transmission. Likewise, the *outgoing* tunnel identifiers and *outgoing* physical interfaces are related to the *sending* step in the claims. There is no reasonable reading of the claims where the sending step is not dealing with an outgoing transmission. The Court should reject IV's alternative suggestion of using "first" and "second" because those terms further obfuscate these elements in the claims. In addition, IV only provides this alternative as to the "tunnel identifier" terms, which would not resolve the ambiguities caused by the "physical interface" terms. Finally, "first" and "second" are not used in the specification to describe these elements. As explained above, the specification consistently uses incoming and outgoing for this purpose.

Ultimately, there is ambiguity in the claims and IV has not identified any problem inherent to using the labels in the specification to clarify which components are being used at which steps. To the contrary, it appears IV prefers to maintain the ambiguity in the claim language. If so, the

Court should hold the claims to be indefinite. Alternatively, the Court should adopt VMware's proposal and use the terms "incoming" and "outgoing" consistently with their use in the specification to clarify the ambiguities in the claim language regarding the physical interfaces and tunnel identifiers used in the various steps of the claim.

D. "customer forwarding [table/information]" ('051 patent claims 1 and 3)

VMware Proposal	IV Proposal
Plain and ordinary meaning	"table(s) containing [a set/sets] of customer specific forwarding information" / "set(s) of customer specific forwarding information"

VMware proposes that no construction of these terms are necessary. The language is not complex and would not be confusing to a jury.

The Court should not adopt IV's proposal. First, IV's proposal for "customer forwarding table" defines the contents of this table broadly (as sets of customer specific forwarding information) without reference to information associating network addresses with outgoing physical interfaces and outgoing tunnel identifiers that is required by other language in the claims. *See, e.g.,* '051 patent at cl. 1 ("the customer forwarding tables associating network addresses with physical interfaces and tunnel identifiers . . .); *see* section IV.C herein (discussing how the concept of "outgoing" interfaces and identifiers relate to the customer forwarding table term).

Second, there is no reason to add IV's proposed language of "specific" or "set/sets" to these claim terms. IV argues that this language is necessary to clarify that there is there is a "single, 'correct' customer forwarding table/information accessible to each customer." Dkt. No. 53 at 44–45. But this argument ignores that there is already language in the claims reciting determination of "the correct" table/information. *See e.g.,* '051 patent at cl. 1 ("determining ***the correct*** customer forwarding table from . . ."); *id.* at cl. 3 ("determining ***the correct*** customer forwarding information from . . .").

Third, it is appropriate for the Court to issue a plain and ordinary meaning construction for this term. The parties apparently do not dispute that the claim language, in the context of the entire claim, require tables (or information) that contain customer specific forwarding information segregated by customer. *See O2 Micro*, 521 F.3d at 1362 (not required to construe limitations if there is no fundamental dispute).

VI. DISPUTED TERMS FROM U.S. PATENT NO. RE44,818 (THE "'818 PATENT")

A. "hierarchical token bucket resource allocation"/ "token" (recited in '818 patent claims 1, 17, 30, 32, 33 and 37–42)

VMware Proposal	IV Proposal
the specific class-based scheduling algorithm known in the art as the "hierarchical token bucket"	Plain and ordinary meaning

IV argues that "there is no need to construe" these terms ("HTB") beyond their plain and ordinary meaning. Dkt. No. 53 at 37. However, resolution of this dispute is necessary because the parties dispute what the plain and ordinary meaning is.

IV's brief (and infringement contentions) confirm that it intends to argue that the plain and ordinary meaning of HTB does not actually mean the HTB algorithm that was well-known to POSITAs, but instead is a "non-limiting" term that could encompass other QoS or scheduling mechanisms. For instance, IV argues that the "claimed invention involves 'a two-tier hierarchical QoS management process . . . employed in a virtual I/O server' in which 'network fabric resources are allocated in a hierarchical arrangement,'" that this "hierarchical resource allocation" can be implemented "using scheduling and queuing methods such as hierarchal token bucket." Dkt. No. 53 at 36. IV's brief repeatedly suggests that the term "HTB" is used in the claims in a "non-limiting" manner. *Id*. at 36–37.[18]

[18] Yet, IV fails to identify the "intended meaning of the term," which is the reason the parties are at an impasse as to the construction of this term. Because IV failed to provide a construction that

But, the claims **are** expressly limited to use of the HTB, and do not extend to any "two-tier hierarchical QoS management process" as IV seems to suggest. This is confirmed in both the claims and the specification. The claims themselves uniformly use the term "hierarchical token bucket" – and not any other QoS or scheduling mechanisms. The specification confirms that the inventors knew that HTB was a **specific one** of many QoS or scheduling mechanisms, and that the inventors intended the term HTB to have its well-known meaning – the specific class-based scheduling algorithm known in the art as the "hierarchical token bucket."[19] The inventors were aware of other QoS and scheduling mechanisms, and chose to limit their claims to just one of them – the HTB. Limiting the claims to HTB in this manner was critical to getting them allowed. Dkt. No. 54-31 at 13 (arguing that HTB was different than other "hierarchical or tree structures for storing resource reservations"). IV should be held to this understanding of the claims.

IV's remaining criticisms of VMware's proposal are not persuasive. IV argues that the articles VMware cites show that "multiple different implementations of a hierarchical token bucket algorithm were known" and that therefore the term "should not be limited to a specific, extrinsic definition as VMware proposes." Not so. Each of these articles show that HTB has one and only one meaning: the specific class-based scheduling algorithm known in the art as the "hierarchical token bucket." It is a proper noun, just like the Statue of Liberty. As Dr. Snoeren explains, and

---

resolves the parties' dispute, the Court should deny enter VMware's proposed construction. *See O2 Micro*, 521 F.3d at 1361.

[19] The specification repeatedly and consistently identifies HTB as a specific QoS mechanism at least 12 times. *See, e.g.* '818 patent at 9:61–65; 9:66–10:2 ("QoS manager 414 … maintain[s] a scheduling mechanism, **such as a HTB** scheduling mechanism, that controls whether packets are forwarded for further processing or enqueued on fabric receive buffer"); 10:15–16 ("Hierarchical token bucket can be considered as a class based scheduling mechanism."); 11:49–50 "using a QoS mechanism **such as hierarchical token bucket** (HTB)"); 17:7–10 ("The SAN QoS manager 424, using mechanism **such as HTB**, determines if the application server 102 associated for the write data has sufficient tokens to transmit the write data to the SANI/O systems.").

explained above, this is how the patent itself uses the term. The extrinsic evidence merely confirms that the HTB was well known in the art, and that the patent itself was referring to that common usage of the term and not anything else.[20] Nor does the article referenced by IV as its Exhibit Q support IV's position. This article refers to HTB as a proper noun consistent with VMware's proposal. This article does not discuss different implementations of HTB, and instead merely discusses how to use HTB in various use cases such as Linux and WLAN.

Regarding the term "token," IV does not criticize VMware's construction or offer its own. Its brief, however, characterizes it as a "standalone" term and suggests IV will argue the term has a meaning that applies more broadly than tokens used in "hierarchical token bucket." *See* Dkt. No. 53 at 37–38. But the term token is not a standalone term – it is used specifically in connection with the HTB. The claims confirm this. For instance, they include steps of "classifying" the packets "relative to the hierarchical token bucket resource allocation to determine a current amount of tokens available," "comparing" the size of the packets "to the current amount of tokens available," and "forwarding" the packets "if the current amount of tokens available are sufficient." *See, e.g.* Claim 42. This confirms that VMware's proposal is correct.

IV has identified extrinsic evidence regarding the term token that suggests it intends to rely on some other meaning for the term. Ex. 47 (Microsoft Computer Dictionary, Fifth Ed., (2002) p. 522) ("A unique structured data object or message that circulates continuously among the nodes of a token ring and describes the current state of the network"); *id*. ("Any nonreducible textual

[20] IV also argues that VMware's ignores the significance of ". . . resource allocation [of bandwidth]" which is a mischaracterization of facts. Dr. Snoeren explains that hierarchical token bucket and token "refer to the very specific method of allocating bandwidth resources referred to in the art as a hierarchical token bucket, or "HTB." *See* Dkt. No. 54-2 (Snoeren Decl.) at 48. Moreover, the claims themselves use "hierarchical token bucket resource allocation of bandwidth" as the antecedent to "hierarchical token bucket resource allocation." *See, e.g.*, '818 patent at 17:51–52, 17:61–62.

element in data that is being parsed, for example, the use in a program or a variable name, a reserved word, or an operator."). This would be contrary to the discussion of token in the claims and specification, and should therefore be rejected.

B. "enforc[e/ing]", "receiv[e/ing]", "classify[ing]", "compar[e/ing]", "forward[ing]", and "buffer[ing]" ('818 patent claims 1, 17, 30, 32, 33, 37, 38, 39, 42)

VMware Proposal	IV Proposal
"enforcing . . . across the physical [storage network] interface of the virtual I/O server" "receiv[e/ing] in the virtual I/O server" "classify[ing] in the virtual I/O server" "compar[e/ing] in the virtual I/O server" "forward[ing] in the virtual I/O server" "buffer[ing] in the virtual I/O server"	Plain and ordinary meaning (for each)

These limitations constitute the HTB steps of the claim (hereinafter, the "HTB Steps"). Beginning with the summary of the invention, the patent makes clear that "the present invention" requires that these limitations be performed **in the virtual I/O server**. '818 patent at 1:66–2:3 ("**the present invention** provides methods and apparatuses directed to managing quality of service (QoS) **in virtual input/output (I/O) servers**." *Id*. at 8:19–21 ("**The present invention** manages QoS of I/O subsystems **in virtual I/O servers** …."). The Federal Circuit has repeatedly held that statements regarding "the present invention" limit the claims and confirm that, in the '818 patent, the claimed steps are required to be performed by the virtual I/O server and not anywhere else. *Honeywell Int'l, Inc. v. ITT Industries, Inc.*, 452 F.3d 1312 (Fed. Cir. 2006) ("We agree with the district court that the claim term 'fuel injection system component' is limited to a fuel filter. . . On at least four occasions, the written description refers to the fuel filter as 'this invention' or 'the

present invention'"); *Edwards Lifesciences LLC v. Cook Inc.*, 582 F.3d. 1322, 1329 (Fed. Cir. 2009) ("Here, the specification frequently describes an 'intraluminal graft' as "the present invention" or "this invention," indicating an intent to limit the invention to intraluminal devices."); *Regents of Univ. of Minnesota v. AGA Med. Corp.*, 717 F.3d 929, 936 (Fed. Cir. 2013).

And in fact, the claims are written from the perspective of the virtual I/O server. The independent claims all include the HTB Steps (enforcing, receiving, classifying, comparing, forwarding, and buffering). The summary of the invention and the entire specification "repeatedly and consistently" confirm that these functions are performed only by the virtual I/O server. Dkt. No. 54 at 54–55. The Federal Circuit has consistently held that where, as here, the patentee has repeatedly and consistently explained its invention, the claims should be so limited. *Microsoft Corp. v. Multi-Tech Systems, Inc.*, 357 F.3d 1347 (Fed. Cir. 2004).[21]

IV argues that VMware's proposals improperly "reads out" preferred embodiments and seeks to "redefine structural relationships" among other claim elements. IV argues that VMware's constructions contradict the specification's supposed disclosure that the "virtual I/O server" is "not necessarily physically distinct from the application server(s) and I/O subsystems that it connects." Dkt. No. 53 at 39. This argument misses the point.[22] As explained above, the "present invention" requires that the HTB Steps (e.g. classifying, buffering, etc.) be performed **by the virtual I/O server** and not the claimed application server or anything else. IV points to nothing in the specification where these functions are performed by anything other than the virtual I/O server.

[21] The Federal Circuit held that "claims must be interpreted in light of the specification" and "repeatedly and consistently" describing a feature of the invention is not "limited to describing a preferred embodiment, but more broadly describes the overall invention." *Id.* at 1347–48.
[22] It is also incorrect. IV misleadingly quotes the "Background" section of the patent. The patent does **not** disclose a virtual I/O server that is physically distinct from the application server.

Not only are VMware's constructions consistent entire specification, they are required by the inventors' clear statements as to what they invented.

C. "maintaining a connection over a network fabric" ('818 patent claims 1, 17, 30, 32 and 42)

VMware Proposal	IV Proposal
"maintaining a connection between the physical interface of the application server and the physical interface of the virtual I/O server over a network fabric"	Plain and ordinary meaning

IV argues that VMware is "attempting to insert additional structural limitations "physical interface of the application server" and "physical interface of the virtual I/O server." Dkt. No. 53 at 40. But, the specification specifically discloses that it is the physical interface of the application server that connects to the network fabric in Figure 2. Dkt. No. 54 at 56. And, the specification repeatedly and consistently refers to the connection with the network fabric as a physical connection. *Id*. ("I/O fabric PHY interface 202 generally refers to the hardware interface or interconnection to the I/O switch fabric . . . Virtual I/O server 106 connects to the I/O switch fabric 104 through I/O fabric interface 110 such as Infiniband ports."). The Federal Circuit has held that "when a patentee uses a claim term throughout the entire patent specification, in a manner consistent with only a single meaning, he has defined that term by implication." *Bell Atlantic Network Services, Inc. v. Covad Comm'cns Grp., Inc.*, 262 F.3d 1258, 1271 (Fed. Cir. 2001). This Court should reject IV's invitation to overturn well-established Federal Circuit precedent and adopt VMware's proposed construction.

IV argues that the "virtual I/O server need not be physically distinct from the application servers and associated virtual network interface(s)" and supports its argument with the following quote from the specification:

> Similarly, **virtual network interface**, in one implementation, emulates an Ethernet NIC. In one implementation, this driver plugs in at the bottom of the

> network stack and provides an Internet Protocol address bridged by the Virtual I/O server 106 onto a LAN.

Dkt. No. 53 at 41; '818 patent at 4:9–13 (emphasis added). But this evidence supports VMware's position. As shown in annotated Figure 2 below, the "virtual network interface" does not connect with the "I/O switch Fabric"; the physical interface of the application server (PHY 202) connects to the I/O switch fabric:

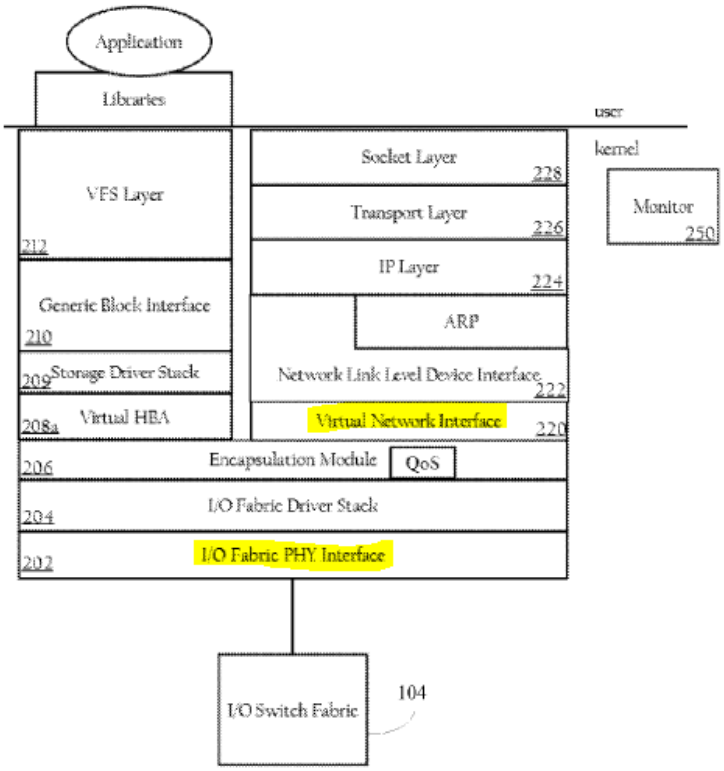

Fig. 2

There is no support in the '818 patent for IV's position that the scope of the claim term can have a broader meaning under the guise of a "plain and ordinary meaning" construction. Especially when the inventor repeatedly and consistently described the invention as maintaining a connection between the physical interface of the application server and the physical interface of the virtual I/O server over a network fabric. The Court should adopt VMware's proposed construction.

D. "virtual storage network interface layer of an application server" / "virtual network interface layer of an application server"/ "virtual interface layer of an application server" ('818 patent claims 1, 17, 30, 32 and 42)

VMware Proposal	IV Proposal
"a virtual storage network interface to higher layers of the virtual node in an application server" / "virtual network interface layer to higher layers of the virtual node in an application server" / "virtual interface layer to higher layers of the virtual node in an application server"	Plain and ordinary meaning

IV identified what it means by "plain and ordinary meaning" for the first time in its opening brief. That is, IV argues that "virtual [network/storage network] interface layer of an application server" means "interface layers (e.g., virtual network interface 220, virtual HBA 208a) that emulate layers of a networking or storage protocol stack" of an application server. Dkt. No. 53 at 35. Given this proposal, VMware proposes the following meaning which is consistent with the intrinsic evidence and IV's proposed construction: "virtual interface layer (e.g., virtual network interface 220, virtual HBA 208a) of an application server that emulates a physical device or a virtual device relative to a physical device." The only dispute between the proposed constructions above is: what does the virtual interface layer emulate? The specification provides the answer:

> The **virtual hostbus adapter (HBA), emulating a physical HBA**, receives SCSI commands for a given device and passes them to the virtual I/O server 106 over the I/O switch fabric. Similarly, **virtual network interface, in one implementation, emulates an Ethernet NIC**.

'818 patent at 4:6–10.

> **HBA module 208a emulates a physical hostbus adapter** relative to the native operating system executed on the application server 102.

Id. at 7:41–43.

> In other implementations, **the virtual HBA module 208 can be configured to emulate a virtual block device relative to the generic block interface**.

Id. at 7:65–67. The specification repeatedly and consistently disclose that the virtual interface layer emulates a physical device (e.g., physical HBA, Ethernet NIC) or a virtual device relative to

a physical device. *See also id*. 6:29–33 ("Virtualization software in the virtual machine server abstracts the underlying hardware by creating an interface to virtual machines, which represent virtualized resources such as processors, physical memory, network connections, and block devices."). Contrary to IV's proposal, there is no support in the intrinsic evidence for the virtual interface layer to "emulate layers of a networking or storage protocol stack." Accordingly, VMware's proposed construction is entirely consistent with the intrinsic evidence and with how a POSITA would understand the meaning of the disputed terms when reading the '818 patent and should be adopted by the Court.

E. "one or more input/output virtualization modules comprising computer-readable instructions operative to cause the one or more processors to" performs functions terms ('818 patent claim 17)

Plaintiff concedes that VMware's proposal should "succeed" if the Court finds that the term "'one or more input/output virtualization modules' is a nonce term that would not be understood by a PHOSITA as denoting structure." Dkt. No. 53 at 43. VMware has made this showing. Under *Williamson v. Citrix* the term "module" is recognized as a nonce word that does not connote structure. 792 F.3d 1339, 1350 (Fed. Cir. 2015). Furthermore, the prefix "input/output virtualization" adds no known structure; nor does the specification describe a structural component for this term. *Williamson,* 792 F.3d at 1351 ("finding means-plus-function when there was "nothing in the specification or prosecution history that might lead us to construe that expression as the name of a sufficiently definite structure"). Nor does the suffix of "comprising computer-readable instructions operative to cause the one or more processors to" add any structural component. *Glob. Equity Mgmt (SA) Pty. Ltd. v. Expedia, Inc.*, No. 2:16-cv-00095-RWS-RSP, 2016 WL 7416132, at *29 (E.D. Tex. Dec. 22, 2016) ("the 'program code for configuring ...' term is governed by § 112, ¶ 6."); *Personal Audio, LLC v. Apple, Inc*., No. 9:09-cv-111, 2011 WL 11757163, at *21 (E.D. Tex. Jan. 31, 2011) ("'processor' cannot describe sufficient structure").

IV's reliance on *Zeroclick, LLC v. Apple Inc*., as support for its position is misplaced. In *ZeroClick*, the court found that the term "user interface code" was not a means plus function term. In so finding, the court found that the term was a "specific reference[] to conventional graphical user interface programs or code, existing in prior art at the time of the inventions." 891 F.3d 1003, 1008 (Fed. Cir. 2018). The Court bolstered its opinion by observing that the specification described conventional graphical user interface code. *Id.* at 1009. In contrast, in the present case, the term "input/output virtualization modules" is not a reference to a conventional structure existing in the prior art, nor does the specification discuss a conventional "input/output virtualization modules" or even recite the term.

IV's remaining arguments are misplaced. For example, IV argues that for the claimed function of "maintain a connection, over a network fabric, to a virtual storage network interface layer of an application server, wherein the virtual storage network interface layer is associated with a virtual storage node identifier" there is sufficient structure recited in the claim of "a network fabric, a virtual storage network interface layer, and a virtual storage node identifier." Dkt. No. 53 at 42. But, IV fails to realize that none of these components perform the "maintain" function. *See Inventio AG v. ThyssenKrupp Elevator Am. Corp*., 649 F.3d 1350, 1356 (Fed. Cir. 2011) (analyzing whether there is "sufficient structure for performing [the claimed] function"). Instead, the "input/output virtualization modules" are recited in the claim as performing the "maintain" function; the components that IV identified are merely the endpoints being connected.

The same rational supports rejecting each of IV's arguments for these terms. For example, the "physical storage network interface" doesn't perform "enforce" function,[23] it is merely the

[23] *See, e.g.,* Dkt. No. 53-4 at 2–3 ("**enforce** a hierarchical token bucket resource allocation of bandwidth **across the physical storage network interface**") (emphasis added).

medium across which the functions are performed. Nor do the presence of other allegedly structural terms dispersed throughout the clams as objects of the functions, or "additional elements external to the disputed elements" (Dkt. No. 53 at 42), bear any relevance to the structure for performing the "maintain", "present", "enforce", etc. functions recited in claim 17—instead, the claim explicitly recites that it is the input/output virtualization module that causes one or more processors to perform these functions.

Finally, IV's alternative identification of structure disclosed in the specification is unsupported. *See* Dkt. No. 53 at 43. IV has made no attempt to show *any* disclosure of these alleged structures performing the claimed functions. Instead, IV merely proffered a lengthy string-cite without substantive analysis. *Id.* (citing to the '818 patent at 2:9–18; 3:9–11; 3:15–30; 3:43–53; 3:60–4:13; 4:27–5:32, 5:55–6:42, 7:26–67; 8:1–12:28; 12:30–13:4; 13:6–14:29; 14:65–15:23;15:46–17:19; Figs. 1–4; and Figs. 11–13). Dkt. No. 53 at 43. Regardless, VMware's expert has reviewed these sections and found that a PHOSITA would not have seen structure in these sections for performing the claimed functions. Dkt. No. 54-2 (Snoeren Decl.) at 52. As such, this claim is indefinite.

VII. CONCLUSION

For the reasons stated herein, VMware respectfully requests the Court adopt its proposed constructions for the disputed terms and phrases.

Dated: March 27, 2020

Respectfully submitted,

WINSTON & STRAWN LLP

/s/ *Katherine Vidal*

Katherine Vidal
Admitted *Pro Hac Vice*
KVidal@winston.com
Michael R. Rueckheim
Texas State Bar No. 24081129
MRuekheim@winston.com
WINSTON & STRAWN LLP
275 Middlefield Road, Suite 205
Menlo Park, CA 94025
Telephone: (650) 858-6500
Facsimile: (650) 858-6559

Thomas M. Melsheimer
Texas State Bar No. 13922550
M. Brett Johnson
Texas State Bar No. 00790975
MBJohnson@winston.com
Michael A. Bittner
MBittner@winston.com
Admitted *Pro Hac Vice*
Winston & Strawn LLP
2121 N. Pearl St., 9th Floor
Dallas, TX 75201
Telephone: (214) 453-6500
Facsimile: (214) 453-6400

Vivek V. Krishnan
Admitted *Pro Hac Vice*
VKrishnan@winston.com
DaWanna McCray
Admitted *Pro Hac Vice*
DMcCray@winston.com
Winston & Strawn LLP
35 W. Wacker Drive
Chicago, IL 60601
Telephone: (312) 558-5600
Facsimile: (312) 558-5700

William M. Logan
Texas State Bar No. 24106214
WLogan@winston.com
Winston & Strawn LLP
1111 Louisiana Street
Houston, TX 77002
Telephone: (713) 651-2766
Facsimile: (713) 651-2700

ATTORNEYS FOR DEFENDANT VMWARE, INC.

CERTIFICATE OF SERVICE

I hereby certify that, on March 27, 2020, the foregoing document was electronically filed with the Clerk of Court using the Court's CM/ECF system which will send notification of such filing to all counsel of record, including counsel of record for Plaintiffs Intellectual Ventures I LLC and Intellectual Ventures II LLC.

/s/ *Katherine Vidal*
Katherine Vidal

IN THE UNITED STATES DISTRICT COURT

WESTERN DISTRICT OF TEXAS
AUSTIN DIVISION

Intellectual Ventures I LLC and Intellectual Ventures II LLC, **Plaintiffs/Counter-Defendants,** **v.** **VMware, Inc.,** **Defendant/Counter-Plaintiff.**	**Civil Action No. 1:19-CV-01075-ADA** **JURY TRIAL DEMANDED**

DEFENDANT VMWARE, INC.'S REPLY CLAIM CONSTRUCTION BRIEF

TABLE OF CONTENTS

TABLE OF AUTHORITIES

Page(s)

Cases

Statutes

TABLE OF ABBREVIATIONS

Abbreviation	Full Name
PTO	United States Patent and Trademark Office
POSITA	Person of Ordinary Skill in the Art
Snoeren Decl.	Declaration of Alex Snoeren, Ph.D. Regarding Claim Construction
'686 patent	U.S. Patent No. RE 44,686
'726 patent	U.S. Patent No. RE 42,726
'937 patent	U.S. Patent No. 6,985,937
'937 FH	File History of U.S. Patent No. 6,985,937
'752 patent	U.S. Patent No. 7,949,752
'051 patent	U.S. Patent No. RE 43,051
'818 patent	U.S. Patent No. RE 44,818

TABLE OF EXHIBITS
FILED WITH VMWARE'S OPENING CLAIM CONSTRUCTION BRIEF

Exhibit	Title	Dkt. No.
Ex. 1	Parties' Agreed Constructions to Claim Terms	54-1
Ex. 2	Declaration of Alex Snoeren, Ph.D. Regarding Claim Construction	54-2
Ex. 3	'937 Patent File History, Applicant Arguments dated November 17, 2003	54-3
Ex. 4	U.S. Patent No. 5,933,603 to Vahalia et al.	54-4
Ex. 5	'937 Patent File History, Non-Final Office Action dated May 10, 2004	54-5
Ex. 6	USPTO Patent Search of claim term "dynamic virtual server mover"	54-6
Ex. 7	Terms in the '726 patent that overlap with disputed claim terms in the '686 patent	54-7
Ex. 8	The parties' proposals for the terms in the '051 involving multiple recitations of "physical interface(s)" and "tunnel identifier(s)"	54-8
Ex. 9	'752 Patent File History, Response to Office Action dated March 16, 2009	54-9
Ex. 10	'752 Patent File History, Office Action dated May 5, 2009	54-10
Ex. 11	'752 Patent File History, Office Action dated October 30, 2009	54-11
Ex. 12	'752 Patent File History, Response to Office Action dated November 13, 2009	54-12
Ex. 13	'752 Patent File History, Office Action dated March 16, 2010	54-13
Ex. 14	'752 Patent File History, Response to Office Action dated August 20, 2010	54-14
Ex. 15	'752 Patent File History, Response to Office Action dated August 5, 2009	54-15
Ex. 16	'752 Patent File History, Final Office Action dated November 8, 2010	54-16
Ex. 17	'752 Patent File History, Response to Final Office Action dated January 4, 2011	54-17
Ex. 18	Report and Recommendation of United States Magistrate Judge, Intellectual Ventures v. HCC Ins. Holdings, Inc., Case No. 6:15-cv-660 (E.D. Tex. August 26, 2016) ("'752 Patent Magistrate Report")	54-18
Ex. 19	Charles Aulds, Linux Apache Web Server Administration, 39 (2001) ("Aulds")	54-19
Ex. 20	Barry Nusbaum, WebSphere Application Servers: Standard and Advanced Features 45 (1999)	54-20
Ex. 21	Ludmila Cherkasova, FLEX: Design and Management Strategy for Scalable Web Hosting Service, 14–15 (Oct. 1999)	54-21
Ex. 22	October 14, 2003 Amendment & Remarks, U.S. Patent Appl. No. 09/526,980	54-22
Ex. 23	Physical Interface, The IEEE Standard Dictionary of Electrical and Electronics Terms (6th ed., 1996)	54-23
Ex. 24	RE43,051 (U.S. Patent Appl. No. 11/858,091) Patent File History, May 17, 2010 Office Action	54-24
Ex. 25	U.S. Patent No. 6,286,047 ("Ramanathan")	54-25

Exhibit	Title	Dkt. No.
Ex. 26	U.S. Patent No. 6,247,057 ("Barrera")	54-26
Ex. 27	Webster definition of consume	54-27
Ex. 28	'752 Patent List of Disputed Claim Terms	54-28
Ex. 29	BEN LAURIE AND PETER LAURIE, APACHE: THE DEFINITIVE GUIDE 163, 177, 242–43, 295 (1999)	54-29
Ex. 30	*Virtual Server*, MICROSOFT COMPUTER DICTIONARY (5th ed. 2002), p. 555 (IV-VMWARE-00004290)	54-30
Ex. 31	'789 Patent (Orig. Patent) prosecution history, 2009.08.25 Resp. to Office Action at 12	54-31
Ex. 32	Competing Parties' Proposals for the '818 Patent	54-32
Ex. 33	Grotto Networking, available at https://www.grotto-networking.com/BBQoS.html	54-33
Ex. 34	http://tldp.org/HOWTO/Traffic-Control-HOWTO/index.html	54-34
Ex. 35	https://lartc.org/howto/lartc.qdisc.classful.html#AEN1071	54-35
Ex. 36	Linux Advanced Routing & Traffic Control HOWTO, located at http://www.oamk.fi/~jukkao/lartc.pdf	54-36
Ex. 37	Traffic Control HOWTO, Version 1.0.2, Martin A. Brown, located at https://www.tldp.org/HOWTO/html_single/Traffic-Control-HOWTO/	54-37
Ex. 38	Traffic Control using tcng and HTB HOWTO, Version 1.0.1, Martin A. Brown, April 2006, located at http://linux-ip.net/articles/Traffic-Control-tcng-HTB-HOWTO.html	54-38
Ex. 39	Bavier, et al., Operating System Support for Planetary-Scale Network Services, Proceedings of the First Symposium on Networked Systems Design and Implementation (March 2004), located at https://www.usenix.org/legacy/events/nsdi04/tech/full_papers/bavier/bavier.pdf	54-39
Ex. 40	Benita, Kernel Korner - Analysis of the HTB Queuing Discipline Software, available at https://www.linuxjournal.com/article/7562 as of January 26, 2005 and printed in the Linux Journal, March 2005.	54-40
Ex. 41	U.S. Patent 7,161,904 titled: "System and method for hierarchical metering in a virtual router based network switch" to Hussein et al.	54-41
Ex. 42	Bavier et al, Container-based Operating SystemVirtualization: AScalable,High-performance Alternative to Hypervisors, Conference Paper in ACM SIGOPS Operating Systems Review, January 2007, located at http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.1018.1012&rep=rep1&type=pdf	54-42
Ex. 43	Valenzuela, J.L., et al., "A Hierarchical Token Bucket Algorithm to Enhance QoS in IEEE 802.11: Proposal, Implementation and Evaluation, IEEE, vol. 4, Sep. 2004 ("Valenzuela Article")	54-43
Ex. 44	Email from J. Deblois to M. Rueckheim dated March 3, 2020	54-44
Ex. 45	Webster's New Work Telecom Dictionary, Definition of Layer	54-45

TABLE OF EXHIBITS
FILED WITH VMWARE'S RESPONSIVE CLAIM CONSTRUCTION BRIEF

Exhibit	Title	Dkt. No.
Ex. 46	U.S. Patent No. 6,351,775 ("Yu")	61-1
Ex. 47	*Token*, Microsoft Computer Dictionary (5th ed. 2002), p. 532 (IV-VMWARE-00004257)	61-2
Ex. 48	U.S. Patent No. 6,976,258 ("Goyal")	61-3
Ex. 49	'937 Patent File History, (January 8, 2004 Office Action)	61-4

TABLE OF EXHIBITS
FILED WITH VMWARE'S REPLY CLAIM CONSTRUCTION BRIEF

Exhibit	Title
Ex. 50	U.S. Patent No. 5,740,231 ("Cohn")
Ex. 51	Supplemental Declaration of Dr. Alex Snoeren
Ex. 52	'937 Patent File History (October 5, 2004 Applicant Arguments)
Ex. 53	Plaintiff's January 31, 2020 Disclosure of Proposed Claim Terms for Construction

I. DISPUTED TERMS FROM U.S. PATENT NO. RE44,686

A. "modif[y/ied] [a] resource allocation" / "modify[ing] [the] computer resources allocated to a virtual server" ('686 patent claims 5–7)

The primary dispute is whether IV can step away from the clear definitional statements and disclaimers in the intrinsic record to expand "modifying a resource allocation" to cover something other than modifying a quality of service guarantee. With respect to the specification, IV fails to address clear definitional statements. *See* '686 patent at 4:49-60 (defining that "[a] resource allocation for a virtual server is specified as a 'quality of service guarantee'"); '686 patent at 2:48-50 ("summary of invention" discussing "the present invention" as "adjust[ing] the quality of service guarantees for virtual servers"). IV cannot rebut the clear impact of these definitions.

With respect to the prosecution history: the applicant was clear when it equated the modifying a resource allocation with modifying a quality of service guarantee. IV does not dispute the record, but takes the tenuous position that the statements are limited to claim 1. IV is wrong.

First, the disclaimer could not have been focused on Claim 1's "requirement that the resource allocation be specified as the claimed quality of service guarantee." Dkt. No. 60 at 9-10. Claim 1 had no such express requirement. The claim did not define any relationship between a "resource allocation" and a "quality of service guarantee." It only recited that a quality of service guarantee was to be increased if it was determined that a process was overloaded.

Second, IV ignores that the applicant made the same disclaimer with regard to the other independent claims, noting that they recite "similar" limitations: "Claims 13 and claims 27 and 37 ***similarly recite*** "increasing the ***resource allocation*** limit of the virtual server"; and claims 39 (as amended) and 44 (as amended) ***similarly recite*** "configured to ... modify a ***resource allocation*** for the virtual server" . . . [and therefore] are also patentable." Dkt. No. 54-3 at 22.[1]

[1] Page citations herein refer to the ECF page numbering unless otherwise noted.

For all claims, and consistent with the specification's definitional statements, the applicant clarified that its general use of the term "resource allocation" was "equivalent" to a "quality of service guarantee." *See* Dkt. No. 54-3 at 20. The applicant specifically equated these two terms to distinguish the prior art that disclosed a static form of resource allocation. *See, e.g.,* Dkt. No. 54-3 at 21 (distinguishing prior art: "[h]owever, the server ***resources allocated to the request do not change***; what does change is which server node handles the request. ***Thus, Yu does not disclose increasing a quality of service guarantee*** for a process, as required by claim 1"). And, indeed the examiner expressly reiterated and underscored the importance of the applicant's disclaimer. *See* Dkt. No. 54-5 at 4-5. The law is clear that once the applicant disclaims claim scope to secure a patent, neither the applicant (nor a successor in interest like IV) can later broaden the scope. Any other result is not supported by the law,[2] would expand patent rights beyond what Congress intended, and would be fundamentally unfair to the public who is entitled to rely on the public record.

Lastly, IV should not be permitted to read the term "allocation" out of the claim. Though the parties agree that the term ***resource*** means a "set of functions and features of a physical host, such as disk space . . .," the specification makes clear that a resource allocation is information that specifies a guaranteed amount of an assigned resource, and not the resource itself. The law is clear that every word in a claim must be given import. *See Pause Tech., LLC v. TiVo, Inc*., 419 F.3d 1326, 1334 (Fed. Cir. 2005) ("must give each claim term the respect that it is due").

[2] *Springs Window Fashions LP v. Novo Industries, L.P.*, 323 F.3d 989, 995 (Fed. Cir. 2003) ("If the applicant mistakenly disclaimed coverage of the claimed invention, then the applicant should have amended the file to reflect the error."); *Fenner Investments, Ltd. v. Cellco Partnership*, 778 F.3d 1320, 1323 (Fed. Cir. 2015) ("Any explanation, elaboration, or qualification presented by the inventor during patent examination is relevant, for the role of claim construction is to "capture the scope of the actual invention" that is disclosed, described, and patented.").

B. "resource unavailable messages resulting from denied requests to modify a resource allocation" ('686 patent claims 5–7)

The primary dispute for this term is again whether IV can read express limitations out of a claim–here "to modify a resource allocation." IV provides no support for its proposal to equate "denied requests to modify a resource allocation" with the general concept of: "resource denials." Dkt. No. 60 at 13-14. Denying a request to modify a resource allocation is different than declining any type of request by a virtual server (as IV describes "resource denials"). IV cannot at this stage rewrite the invention.[3]

There are numerous problems with IV's proposal. For example, IV's proposal does not include its proposed construction of "modify a resource allocation", even though this language is in the disputed term. *See, e.g.,* Dkt. No. 61 at 14. IV's proposal also injects indefiniteness into the claims through the phrase "cannot be immediately serviced." IV provides no objective standard to measure the scope of the term "immediately." *See, e.g., infra*, Section II(A).[4] Furthermore, IV attempts to change plural nouns—"denied ***requests***" and "resource unavailable ***messages***"—to singular nouns—"a request" and "an indication," respectively.

IV's purported support for its position is either inapposite or unsupported attorney argument. IV relies upon the specification at 2:62-65 which describes "determin[ing] if a virtual server is overloaded by monitoring resource denials." This section does not discuss an embodiment of "resource unavailable messages ***resulting from denied requests to modify a resource allocation***" and is thus inapposite. IV's argument that "[a] denied request to modify a resource allocation, alone is not enough to cause the computer system to take action" because "an instruction must cause the system to do so" is unsupported attorney argument.

---

[3] As to the language "resource unavailable messages," VMware and IV's proposals overlap with the arguments for the disputed '726 patent term "resource denial." *See infra* Section II.A.

[4] VMware submits that IV's proposal renders the claim indefinite under 35 U.S.C. § 112.

C. "determination that a virtual server is overloaded" ('686 patent claims 5–7)

See infra Section II(C), discussing the '726 patent term "virtual server overloaded signal."

D. "virtual server" ('686 patent claims 5–7)[5]

The primary dispute is whether IV should be rightly bound by a definitional statement provided in the specification which is consistent with how a POSITA would understand the term.

First, IV tries to use unsupported expert testimony parroting attorney argument to distance itself from the specification. Dkt. No. 60 at 17; Dkt. No. 60-1 at ¶ 31. The law is clear that (1) extrinsic evidence cannot override the clear intrinsic record;[6] and (2) unsupported, conclusory expert testimony should be given little if any weight.[7] Further, IV's expert has not identified any disagreement with VMware's expert's detailed analysis of this term. VMware's unrebutted and substantive expert opinion as to how a POSITA would have understood this term consistent with the speciation compels the adoption of VMware's proposal. Dkt. No. 54-2 at ¶¶ 34, 90-101.

Second, IV argues that the prosecution history of the parent '937 patent supports its position. Not so. VMware's responsive brief shows that the examiner viewed "virtual server" as a type of "process," consistent with VMware's proposal. See Dkt. No. 61 at 17-18.

Nor does the amendment changing "process" to "virtual server" support IV's position. This amendment was a word substitution without ***any*** argument that "process" and "virtual server" mean different things. Indeed, both before and after this amendment the applicant did not dispute the applicability of the prior art to these terms. *Compare* Dkt. No. 54-3 (2003-11-17 Applicant

---

[5] "Virtual server" is also present in the '051 patent, whose parent application was originally assigned to Ensim and filed nearly concurrently with the parent application of the '686 patent.

[6] *Vitronics Corp. v. Conceptronic, Inc*., 90 F.3d 1576, 1583 (Fed. Cir. 1996) (if "intrinsic evidence alone will resolve any ambiguity . . . it is improper to rely on extrinsic evidence.")

[7] *Phillips v. AWH Corp*., 415 F.3d 1303, 1318 (Fed. Cir. 2005) ("conclusory, unsupported assertions by experts as to the definition of a claim term are not useful to a court").

Arguments) and Ex. 52 (2004-10-05 Applicant Arguments). Thus, this amendment simply indicates that both examiner and applicant understood virtual server to be a type of process, as VMware proposes.

Third, IV argues, without support, that a process cannot provide the "high server availability and large amounts of processing power" discussed in the patent. Dkt. No. 60 at 18-19. Rather, it is common knowledge that different processes, depending upon their complexity, require different amounts of processing power.

Fourth, IV argues that VMware's proposal is based on disclosure of a single embodiment. Not so. VMware derives support from a definitional statement that applies to the entire patent. '686 patent at 3:53-55 ("The term 'virtual server' ***as used herein*** refers to a virtual server capable of receiving a quality of service guarantee from a physical host.") (emphasis added).

E. "determining that a second physical host can accommodate the requested modified resource allocation" ('686 patent claims 5–7)

IV's brief makes this claim term no less indefinite. Specifically, IV argues that the ***singular*** term "the requested modified resource allocation" is referring back to the ***plural*** term "denied requests to modify a resource allocation." The intrinsic record offers no guidance to a POSITA on interpreting this term. IV's post-hoc explanation cannot cure indefiniteness.

IV cites to the '686 specifications at 11:12-57, but this citation is not helpful. This section describes using an easiest fit heuristic model to determine "the physical host that has the most available resources." '686 patent at 11:27-55. But determining which host has the most resources is not the same as the claim requirement of determining that a host "can accommodate ***the*** requested modified resource allocation" where there are multiple "denied requests to modify a

resource allocation." IV has not shown how a POSITA would know which of the denied ***requests*** to modify a resource allocation are used for determining the capacity of the second physical host.[8]

F. "component configured to" Means-Plus-Function Terms ('686 patent claim 7)[9]

As a preliminary issue, IV's arguments for these terms are based on a fundamental misreading of the claim language. Recognizing it cannot avoid means-plus-function treatment without importing structure into the claim, IV argues that the "component configured to" terms are recited as ***within a physical host*** that includes one or more processors and one or more memories." Dkt. No. 60 at 24 (emphasis added). Not true. The claim actually recites (i) "a virtual server operating ***in a first physical host***"; (ii) a ***separate*** "system for modifying the computer resources allocated" to the virtual server; and (iii) the claimed "component" terms as parts ***of the system***. There is no recital of the "components" being part of physical hosts or having any structure.

a. "a component configured to receive an indication . . ." term (see Dkt. No. 53-2 clause 1)[5]

IV's argument is circular and unsupported. IV argues that the "load balancing calculator 530" is the "structure that ***receives*** the indication that a first physical host is overloaded." Dkt. No. 60 at 25 (emphasis added). But this argument ignores that the calculator is described as the component that "***determines*** that physical host 160A is overloaded." '686 patent at 11:24-26 (emphasis added). It defies common sense that one structure will both determine an overload while also receiving an indication of overload.

Nor is there description of the calculator receiving the claimed indication. The calculator determines overload based on analyzing whether the host would support a resource allocation increase request. *See* '686 patent at 11:4-26. The component of the claim, in contrast, ***receives*** that

[8] IV cites to the expert declaration of Dr. Akl for support. But, this declaration simply parrots IV's attorney arguments verbatim without providing any technical analysis whatsoever.
[9] The parties proposed terms and construction are identified at Dkt. No. 54 at 21–24.

indication "based on ***a determination that a virtual server is overloaded*** and wherein the determination that a virtual server is overloaded is ***based on one or more resource unavailable messages resulting from denied requests to modify a resource allocation***." '686 at claim 7 (emphasis added). Regardless, the calculator component is a black box term stated in purely functional terms (load-balancing) and does not connote structure.

b. "component configured to determine that a second physical host can accommodate the requested modified resource allocation" term (see Dkt. No. 53-2 clause 2)

IV again points to an alleged structure that does not perform the claimed function. IV argues that the load balancing calculator's use of the "easiest fit heuristic" is the structure. Dkt. No. 60 at 25-26 and n 10. But this argument ignores that the "easiest fit heuristic" merely determines "the physical host that has the most available resources." '686 patent at 11:27-57. This heuristic does not determine whether a host can accommodate a specific "requested modified resource allocation" as claimed. *See, e.g., supra* at Section I(E).

c. "component configured to generate a physical host transfer signal that indicates a second physical host and to transfer the virtual server . . ." term (see Dkt. No. 53-2 clause 3)

IV agrees that the structure includes at least VMware's proposal of structure, while arguing that "VMware's structure is too narrow." Dkt. No. 60 at 26. Therefore, at a minimum, the construction of this term should include VMware's proposed structure of "Dynamic Virtual Server Mover 140 as described in the '686 patent, 12:1-28; Figure 6." To minimize the disputes, VMware is willing to agree to adding IV's additional structures of: "load balancing module 130," "load balancing calculator 530," "virtual server resource monitor 120," and "physical host resource monitor 540" to the extent structural descriptions are provided of these terms in the specification.[10]

[10] IV identified these alleged structures for the first time in its responsive claim construction brief.

II. DISPUTED TERMS FROM U.S. PATENT NO. RE42,726

A. "resource denials" ('726 patent claims 1, 4–5 and 8)

The primary dispute for this term is whether permissive language in the specification can trump clear definitional statements. IV relies on permissive language in the '726 patent at 2:55-61 which states "***[i]n one embodiment***, . . . [a] resource denial ***may refer*** to any request by the virtual server that cannot be immediately serviced" (emphasis added). In contrast, VMware's proposal derives from a definitional statement: "resource denials, ***which are*** instances wherein a request for additional resources is either implicitly or explicitly denied." '726 at 7:58-60 (emphasis added). IV should be bound by the definition which gives clear public notice, and should not be permitted to rewrite the claims by importing permissive language relating to a single embodiment.

Beyond that, IV's proposal is again unsupported by the plain language of the claims and would render the claim term indefinite under 35 U.S.C. § 112.[11] IV's proposal reads out the claim term "resource," and thus, improperly expands the term to encompass denied requests for something other than additional resources. Additionally, IV's proposal again imports a vague and unfounded "immediacy" requirement replacing the claim term "denials" with "cannot be immediately serviced." IV does not and cannot provide an objective standard to measure the scope of the term "immediately." VMware's proposal, in contrast, is consistent with the plain claim language and the specification. *See, e.g.,* 5:28-32 and 7:54-61 (describing requests for additional resources such as disk space and memory being denied when the resource is fully used).

B. "quality of service guarantee" ('726 patent claims 1 and 4)

IV's arguments for this term are inconsistent and hard to follow. For example, IV argues that "the plain language of IV's proposal specify[ies] that an '***amount of an assigned resource***' is

[11] Should the Court adopt IV's proposal, it should likewise find the claim term indefinite.

the item that is guaranteed." Dkt. No. 60 at 22 (emphasis added). But that language is in VMware's proposal, not in IV's. IV then inconsistently argues that VMware's inclusion of this identical language ("***assigned resource***") somehow "interject[s] ambiguity" into the claim. *Compare* Dkt. No. 60 at 22 with *id.* at 22-23 (emphasis added). This is nonsensical. IV also argues that the term "assigned resource" somehow makes the quality of service guarantee "fixed." But VMware's proposal does not at all suggest that the guaranteed amount of assigned resource is static; indeed, VMware's proposal expressly includes: "that can be dynamically increased/modified."

Similarly, IV argues that "'[q]uality of service guarantee information' is described in the specification as distinct from a 'quality of service guarantee'." Dkt. No. 60 at 23 (citing to the '726 patent at 4:39-60). But, IV cites to a portion of the specification that expressly equates the terms:

> A resource allocation for a virtual server is specified as a 'quality of service guarantee' for that particular server. Each physical host ***stores quality of service guarantees*** for the virtual servers it hosts. As a physical host performs processes associated with a particular virtual server, the physical host ***accesses the stored quality of service information*** to enable the physical host to request the correct quality of service from the operating system kernel of the physical host.

'726 patent at 4:39-60 (emphasis added).

C. "virtual server overloaded signal" ('726 patent claims 1, 4-5 and 8)

The primary dispute is whether IV should be bound by statements made in the specification. VMware's proposal is grounded in the written description which only discloses one way to determine that a virtual server is overloaded. That disclosure is at '686 patent column 5:42–50, which discloses a determination is based on "an average number of resource denials" being "beyond a pre-configured threshold." *See also* '686 patent at 9:4-10; 9:22-28; 9:42-44.[12]

[12] The '726 patent shares an identical specification.

IV, without support, proposes "an indication that a virtual server has been or is being denied resources." Dkt. No. 53 at 25. While IV identifies the '686 patent at 2:62-65, 5:15-17 and 8:5-55 (Dkt. No. 60 at 16),[13] the 2:62-65 and 5:15-17 sections merely use the term "overloaded" at a high level, without disclosing what it means for a virtual server to be overloaded. Similarly, the 8:5-55 section discusses system calls without disclosing an overloaded determination. IV's citation to the '686 patent at column 2:62-65 also ignores that this paragraph concludes in accordance with VMware's proposal by stating: "[i]f the resource denials received by a particular virtual server exceed a pre-specified limit, the virtual server is considered overloaded." '686 patent at 3:1-3.

D. Mean-Plus-Function Elements ('726 Patent claims 1, 3, 4, 5, 7) (see Dkt. No. 53-2 Clauses 4-8)

IV recognizes that structure is needed in the claim to avoid a means-plus-function interpretation, but is also wrong that each of these terms is recited as "tak[ing] place in a physical host server." Dkt. No. 60 at 27. The actual claim language recites (i) a virtual server operating in a ***physical host***; (ii) a ***separate*** "system for dynamically modifying the computer resources allocated to a virtual server; and (iii) the claimed terms at issue (monitor, modifier, module, and mover) as part ***of the system***. There is no recital of these terms being part of physical hosts.

As to the structural disclosure provided by the specification, IV admits that VMware's proposed structure is correct, with the argued caveat that the structure is too narrow. *See* Dkt. No. 60 at 28-30. Therefore, at a minimum, the construction of these terms should include VMware's proposed structure. IV's attempts to broaden these structure, though, are unsupported attorney arguments. For example, IV cites to its expert declaration, but that declaration provides no

[13] The parties dispute the meaning of the '686 patent's similar term, "determination that a virtual server is overloaded". Although, IV has inconsistently proposed a construction for the '726 term and not for the '686 term. VMware proposes the same construction for both patents.

substantive analysis and merely parrots IV's attorney arguments. Also, IV's argument that a POSITA may know other, undisclosed, structure misunderstands the law.[14] Under 35 U.S.C. § 112, ¶ 6, the Court is required to look to the "corresponding structure, material, or acts described in the specification and equivalents thereof." Undisclosed structure is not enough. *See Blackboard, Inc. v. Desire2Learn, Inc*., 574 F.3d 1371, 1382 (Fed. Cir. 2009).

Finally, IV is wrong to argue that VMware is attempting to "read the term 'creating' into each one of its proposed functions." Dkt. No. 60 at 27. VMware's proposal includes brackets to specify that "creating" only applies for claim 4's use of the word. *See* Dkt. No. 60 at 27.

III. DISPUTED TERMS FROM U.S. PATENT NO. 7,949,752

A. "exhausted" ('752 patent claims 1, 9 and 24)

The dispute is whether IV should be bound by a disclaimer repeatedly made during prosecution. As a preliminary issue, IV focuses on the prosecution discussion of the Chou and Cohn references. But IV failed to address, or even discuss, applicant's disclaimer that the Humpleman prior art's disclosure of ***reusable hardware*** does not meet the claim language of a service resource that is consumed and exhausted. *See* Dkt. No. 54 at 35. Indeed, though not required for a disclaimer to have effect,[15] the applicant was successful with this argument. In the very next office action, the examiner dropped the Humpleman reference and applied the Chou reference which disclosed a ***software*** service resource. *Id.* IV should be held to this disclaimer. *See Springs Window*, 323 F.3d at 995 ("The public notice function of a patent and its prosecution history requires that a patentee be held to what he declares during the prosecution of his patent.").

[14] *See, e.g.,* Dkt. No. 60 at 28 (citing to '726 patent 8:46-49; 8:65-66; 9:43-46—each of which are inapposite or merely state there may be other, undisclosed, ways known); *see also id.* at 30-31 (arguing that the "the 'make, then break' and 'break, then make' methods of transfer are not the only methods for transferring virtual servers and thus the structure should not be so limited").
[15] *See Laitram Corp. v. Morehouse Indus., Inc*., 143 F.3d 1456 (Fed.Cir.1998).

The unrebutted Humpleman disclaimer is dispositive. "A patentee may not state during prosecution that the claims do not cover a particular device and then change position and later sue a party who makes that same device for infringement." *Id.* Yet, that is exactly what IV is doing here. *See, e.g.,* Dkt. No. 9 at ¶¶ 47-53 (accusing VMware's use of reusable hardware, such as CPU and memory, of infringement). This is also what IV is arguing in its responsive brief. *See* Dkt. No. 60 at 32 n 13 (arguing that service resources include reusable hardware such as "phone lines, switches, etc."). This attempt to recapture disclaimed subject matter is improper.

The public notice provided by applicant's repeated, consistent and un-abandoned arguments during the entirety of prosecution is also dispositive. IV first narrowed the claim scope from merely ***using*** a service resource to being configured to be ***consumed and exhausted***. *See* Dkt. No. 54 at 35. IV then argued ***successfully*** against the reusable hardware disclosure in Humpleman as not being consumed and exhausted. *Id*. (showing examiner dropped the Humpleman reference in favor of the Chou software reference). IV then argued ***successfully*** that the reusable software in Chou ***is the exact opposite*** of the claimed service resource configured to be consumed and exhausted. *Id.* at 35-36 (showing how the examiner dropped the software Chou reference in favor of the Cohn non-reusable fees reference).[16] IV then distinguished the final Cohn reference by doing two things: (i) amending the claims to include a URL requirement and (ii) "travers[ing]" the prior rejection "[w]ithout acquiescing to the propriety of the rejection." *Id.* at 54. This prosecution history paints a single picture: a non-abandoned disclaimer of claim scope.

[16] IV's argument that the Chou rejection was dropped because of the "an amount of" claim amendment is unsupported speculation. Dkt. No. 60 at 31. This argument is also shown false by the fact that the applicant's disclaimer used ellipses to omit the "an amount of" language. Dkt. No. 54-14 at 10 ("Rather in contrast to "[a] . . . resource [that] is ***exhausted*** upon being ***consumed***," as recited, using respective language, by claims 86, 92, 107, and 126, Chou teaches the exact opposite – ie that the "cartridges" in Chou are ***reused***.") (emphasis and ellipses in original).

IV's attempt to muddy the record with respect to the Chou disclaimer does not warrant allowing IV to recapture disclaimed claim scope. For example, IV focuses on examiner's comment that "[o]n page 25 of the applicant's specification resources are defined as including processing time, memory storage space, and the like . . . [which] can always be reused later." *See* Dkt. No. 60 at 30-31. But IV's speculation aside, the examiner appears to merely be clarifying types of resources because the comment is directed to the description of ***computing resources*** in the specification rather than ***service resources***. *See* '752 patent at 8:19-24. Regardless, with respect to the scales of public notice: this examiner statement does not outweigh the applicant's prior repeated, and later maintained, arguments that prior art disclosures of reusable hardware and software are not within the scope of the invention.[17]

IV's characterization of the Cohn references is incorrect. Cohn did not disclose reusable service resources such as "information providers". Dkt. No. 60 at 31. The examiner explicitly stated that the limitation of "an amount of the service resource is exhausted upon being consumed by the agent," is met by Cohn's disclosure that "some amount is clearly consumed and the user is charged for that amount." Dkt. No. 54-16 at 5. In this way, the examiner equated the claim term to consumption/exhaustion of a user specific non-reusable resource.[18] Indeed, the claims

[17] Moreover, IV's focus on the characterization of the applicant's argument as "misleading" misses the point. The applicant never explicitly rescinded this argument. *See, e.g., Uship Intellectual Properties, LLC v. United States*, 714 F.3d 1311, 1315-16 (Fed. Cir. 2013) ("The analysis focuses on what the applicant said, not on whether the representation was necessary or persuasive . . . ."); *Hakim v. Cannon Avent Group PLC*, 479 F.3d 1313, 1317-18 (Fed. Cir. 2007) ("Although a disclaimer made during prosecution can be rescinded, permitting recapture of the disclaimed scope, the prosecution history must be sufficiently clear to inform the examiner that the previous disclaimer, and the prior art that it was made to avoid, may need to be revisited.").

[18] *See also* Ex. 50 (U.S. Patent No. 5,740,231 ("Cohn")) at 34:11-13 ("An account is maintained ***for each user*** and the costs of services provided by communications system 10 are determined by each user's class of service") (emphasis added).

themselves are similarly limited to user-specific consumption/exhaustion of service resources as shown by their recital of "user customized network-based operations" where a service resource is exhausted upon being consumed by the ***user-specific*** "network-based agent". *See, e.g.*, '752 patent at claim 9.[19] Therefore, VMware's proposed construction should be adopted.

B. "consumed" (recited in '752 patent claims 1, 9 and 24)

The primary issue for this term is whether IV is bound by definitional statements provided in the specification and prosecution history. With respect to the specification: IV argues that the Court should not adopt the definitional statement that VMware's proposal relies upon because the term is "referred to interchangeably with 'used' at least half a dozen times." Not so. The specification explicitly defines "consumed" as "used up" using quotation marks which is a "strong indication" of a definition being provided. '752 patent at 8:22; *Intellectual Ventures I, LLC v. Lenovo Group Ltd.*, 365 F.Supp.3d 200, 206 (D. Mass. 2019).

With respect to the prosecution history: IV has failed to explain how "consumed" can mean "used" (as IV proposes) when the applicant first tried to broadly claim an "agent, for ***using*** a service and a service resource," but later narrowed the claims to recite something different: an "agent, for using a service and service resource ***configured to be consumed***" to overcome prior art. *See* Dkt. No 54 at 38-39. IV argues that this claim amendment was "made to specify that the agent is doing the consuming." Dkt. No. 60 at 33. But IV fails to acknowledge that claim language prior to the amendment already recited that the agent uses the resource. Dkt. No. 61 at 36.

IV is also incorrect in arguing that VMware's proposal is a "transparent attempt to save its carefully crafted non-infringement position regarding the 'exhausted' term." VMware proposed

[19] The claims' recital of user-specific consumption/exhaustion of services resources also belies IV's argument referring to a scenario where "subscriber A gets access to the database for information Y" and "information Y is not then 'unavailable for reuse,' [because] subscriber B can also subscribe to and receive access to information Y." Dkt. No. 60 at 31-32.

that the Court construe the "exhausted" term as that term is dispositive. But, VMware only proposed a construction for "consumed" after IV identified it as needing construction. Ex. 53 (2020-01-31 Plaintiff's Disclosure of Proposed Claim Terms for Construction) at 2.

C. "service" ('752 patent claims 1, 3, 9 and 24)

The primary issue for this term is whether the specification better supports VMware's or IV's proposal. VMware's proposal should be adopted. The specification consistently discloses that the "service" is an application. *See, e.g.*, '752 patent at 3:3-5 and 10:17-21.

IV's arguments do not warrant a different result. For example, IV's argument that the "specification notes that 'an agent is operable to use a service'" (Dkt. No. 60 at 34) does not speak to whether the service is an application or something else (undisclosed). Moreover, the claims already recite that an "agent" is operable "for using a service." Dkt. No. 61 at 37. Additionally, IV's argument that the specification's disclosure of "email service, a voice mail service, etc." are "network functionality, not applications" (Dkt. No. 60 at 34) ignores that these disclosures are in fact applications (e.g., in modern parlance: Microsoft Outlook) that are performed "on behalf of [a] respective principal" as VMware proposes. *See, e.g.,* '752 patent at 8:58-64.

D. Means-Plus-Function Terms

The primary issue for these terms is whether the Court's construction should identify where the specification discloses the structure for these means-plus-function terms. IV's only argument is that the "Court has previously construed means plus function terms where the corresponding structure for performing the claimed function includes only the disclosed structure, e.g., 'widget 26,' and § 112 ¶ 6 requires nothing more." Dkt. No. 60 at 35 (citing to *VLSI Tech. LLC, v. Intel Corp.*, Case No. 1:19-CV-00977-ADA (W.D. Tex.), Claim Construction Order, Dkt. No. 101). Not true. The Court in Case No. 1:19-CV-00977 did not construe a structure of "widget 26". Instead, the Court construed a Section 112 ¶ 6 term in terms of where it is disclosed in the

specification as VMware proposes. *See, e.g.*, *id.* at 2 (construing "Structure: '(1) Voltage source VREF coupled in series with conductor 37, ***as shown in Figure 2***"') (emphasis added).

With respect to the "means for monitoring" term, IV argues that "VMware's proposal includes only a sub-component of the service wrapper 26 which performs the function as a whole." Dkt. No. 60 at 35. Not so. The service wrapper 26 includes a monitor (which VMware propose as the disclosed structure for performing the claimed function) and a converter. IV argues that the converter is part of the structure for performing the claimed function because "the converter converts the computer language used by a service to the computer language used by the agent server" and "in turn interacts with monitor 50 to enforce service resource usage of the agents when using the service." *Id.* But this argument ignores that the claimed "monitoring" is clearly performed by the "monitor" as referenced in the '752 patent at 16:50-61. That other components, such as the "converter", the "agent", the "agent server", and/or the "communication line" are also involved in interacting with the monitor is irrelevant. Under Section 112 ¶ 6 the Court is required to look for the structure described in the specification as performing the claimed function and not ancillary structures disclosed as interacting with relevant structure.

IV. DISPUTED TERMS FROM U.S. PATENT NO. RE43,051

A. "virtual server" ('051 patent claims 1, 3 and 6)

The primary issues for this term are whether the term should be construed, and if so, should the construction refer to a "process" or a "virtual machine". With respect to the first issue, *O2 Micro* requires a construction where, as here, the parties dispute the ordinary meaning. *O2 Micro Int'l Ltd. v. Beyond Innovation Tech. Co.*, 521 F.3d 1351, 1361 (Fed. Cir. 2008).

With respect to the second issue, VMware contends, as supported by the intrinsic record, that a POSITA would have understood the term to refer to a process. Dkt. No. 54 at 43-44. IV, based solely on extrinsic evidence, attempts to redraft the term as a virtual machine. VMware is

correct. Dr. Snoeren explains that "POSITA would have understood this claim term to refer to a process executing on a host computer that accepts communications requests." Dkt. No. 54-2 at 42-47. IV's expert, in contrast, argues, circularly, that term "'virtual server' is referring to a virtualized server" without substantive expert analysis. Dkt. No. 60-2 at 11.

The specification confirms VMware's proposal. The '051 patent explains how to create the virtual servers by citing an incorporated reference that issued as U.S. Patent No. 6,976,258. '051 patent at 4:64–67; *Meetrix IP, LLC v. Citrix Sys., Inc.*, No. 1:16-CV-1033-LY, 2017 WL 5986191, at *6 (W.D. Tex. Dec. 1, 2017) ("Documents incorporated by reference are intrinsic evidence for purposes of claim construction."). The '258 patent describes a "server" as "a process, executing on a dedicated physical services client [that] services client requests for a single network address (physical host) only." Dkt. No. 61-3 at 13 (3:23–25). These servers can provide *virtual server* functionality through virtual hosting (*i.e.*, servicing requests for multiple network addresses on a single physical host) by "creat[ing] child processes to service the requests." '051 patent at 1:24–44. As this indicates, the *virtual servers* of the invention refer to the functionality of process-based servers to service requests on a single physical host for multiple network addresses.

IV is mistaken that the processes described in the written description are referring to the ***physical*** server and not the ***virtual*** server. The '051 patent makes this distinction itself, noting that a "service provider may utilize ***one physical host computer*** to provide commercial host services to multiple customers" by using "***a server application*** executing on a single physical host [that] can be programmed to process requests made to multiple network addresses." '051 patent at 2:47–53 (emphasis added). Put differently, the physical host computer is distinct from the processes running on it, and the processes are what provide the functionality of the virtual server.

It is also not true that VMware's proposal fails to address the invention's virtual aspect. The '258 patent describes how a process provides for virtual hosting (*i.e.*, servicing requests for multiple network addresses on a single physical host) by "creat[ing] child processes to service the requests," Dkt. No. 61-3 at 12 (1:24–44). This ability for a single application to create child processes to handle client requests on a plurality of network addresses is what provides for the *virtual* aspect of the virtual servers, because the application can function as multiple servers. *See id.* at 1:45–2:5; Dkt. No. 54-2 at 43–44. The Court should therefore adopt VMware's proposal.

B. "physical interface[s]" ('051 patent claims 1 and 3)

The fundamental dispute is whether the plain and ordinary meaning of "physical interface" requires hardware. IV does not address this dispute. The Court should construe this term as VMware proposes because the parties dispute the ordinary meaning. *O2 Micro* 521 F.3d at 1361. IV's argument that the physical interface does not provide a point of communication between two or more devices is a distraction, and it is wrong. *See, e.g.,* Dkt. No. 61 at 43.

C. physical interfaces and tunnel identifiers in the storing / receiving / determining / sending terms ('051 patent claims 1 and 3)

The fundamental dispute for these terms is whether it would help a jury to identify which of the several tunnel identifiers and physical interfaces introduced in the claims are used at each step. VMware's proposal would clarify the claims by indicating which elements are being used at each step. IV's proposal would leave the confusion in place.

Moreover, IV is wrong to side-step the antecedent basis problem. Claim 1, for example, requires "receiving, over a tunnel, a transmission on ***a physical interface***, the transmission containing ***a tunnel identifier***." Later, claim 1 requires "determining via the customer forwarding table ***a physical interface and tunnel identifier***." Claim 1 also involves storing tunnel identifiers and physical interfaces in two distinct locations, customer lookup tables and customer forwarding

tables. The antecedent basis problem arises because it is unclear from the claims which of these numerous physical interfaces and tunnel identifiers are being used at the other steps.

VMware agrees with IV that some instances of physical interface and tunnel identifier are associated with the customer lookup table / information, while other instances of those terms are associated with the customer forwarding tables / information. This is precisely why VMware's proposal is necessary, because it will clarify these associations for the jury.

IV is also wrong that identifying some physical interfaces and tunnel identifiers as incoming and others as outgoing would improperly limit the claims to a preferred embodiment. The plain language of the claims makes this distinction, and the specification supports it. The *incoming* tunnel identifiers and the *incoming* physical interfaces are associated with the *receiving* step in the claims, which necessarily deals with *incoming* transmissions. The *outgoing* tunnel identifiers and *outgoing* physical interfaces are related to the *sending* step in the claims, which necessarily deals with *outgoing* transmissions. This is also true of the specification, which uses the words *incoming* and *outgoing* to make this same distinction throughout. *See id.* at 11:30–32, 11:64–12:1, 12:10–20, 12:30–33, 12:59–63, 13:5–11 Figs. 8–9.

The Court should reject IV's argument that the "specification contemplates numerous instances where the claimed tunnel switching activity can operate in either or both directions." Dkt. No. 60 at 41. VMware's proposal does not speak to whether the tunnel switching activity can operate in either direction. Instead, VMware's proposal simply clarifies the antecedent basis problem created by the claim's recital of multiple "a physical interface", etc. terminology.

D. "customer forwarding [table/information]" ('051 patent claims 1 and 3)

The crux of the dispute on these terms is whether the meaning of customer forwarding table (or customer forwarding information) is unclear. VMware proposes that the Court need not

construe these terms separately. There is no ambiguity for these terms, provided that the antecedent basis issues identified in Section IV(C) are addressed.

IV's proposal seeks to inject the words "set/sets" without justification. For example, IV argues that it seeks "to help one differentiate between singular versus plural instances of the disputed claim terms, and further to maintain consistency between claims 1 and 3." Dkt. No. 60 at 42. But there is no need for clarification. The parties appear to agree that the claim language requires tables (or information) that contain customer-specific forwarding information, segregated by customer. And it is clear from the claim that the customer-specific forwarding information must associate network addresses with outgoing physical interfaces and outgoing tunnel identifiers. There appears to be no reason to, as IV proposes, add additional words to these terms.

V. DISPUTED TERMS FROM U.S. PATENT NO. RE44,818

A. "hierarchical token bucket resource allocation"/ "token" (recited in '818 patent claims 1, 17, 30, 32-33 and 37–42)

The dispute here is a simple one. VMware contends that the claim term "hierarchical token bucket resource allocation" should mean just that – the specific scheduling algorithm known in the art by that very name. IV argues for plain and ordinary meaning without explaining why VMware's proposal is ***not*** the plain and ordinary meaning or offering an alternative definition. As explained in VMware's prior briefing, the inventors knew there were many available techniques for scheduling network traffic but chose to limit their claims to just one: hierarchical token bucket. They even argued for the patentability on that very basis. Thus, VMware's construction is not just the plain and ordinary meaning of the term, but is the only reasonable meaning of the term.

IV's response actually confirms that VMware's proposal is correct. First, IV argues that "the claim drafters chose to expressly recite the term as using the indefinite article 'a'." Dkt. No. 60 at 44. Not so. The claim drafters were required to use an indefinite article because of the

antecedent basis rule. *See* MPEP § 2173.05(e). IV then argues that because HTB is an acronym, it is not a pronoun. This is a red-herring, HTB is used as an acronym *and* a pronoun throughout the specification. *See, e.g.*, '818 patent at 14:23 ("the HTB"); *id*. at 11:7 and *id.* 14:23 (same).

Next, IV argues that because "numerous implementations of a hierarchical token bucket algorithm have been introduced into the art," VMware's proposal should be rejected. Dkt. No. 60 at 44-45. This argument is unsupportable. Just because HTB can be implemented in different ways does not change the fact that HTB is a specific, well-known algorithm with a specific, well-known meaning. While there may be many ways to implement HTB, the term carries a consistent meaning in the art. The essential elements of the HTB are captured at 10:15-29, 31-36 of the '818 patent. As Dr. Snoeren explains, this portion of the patent confirms that the inventors understanding of HTB was consistent with the use of the term in the art. Ex. 51 (Supp. Snoeren Decl.) at ¶ 7. And, the essential elements of the HTB are reflected consistently in the patent and across the art. *Id*. Federal Circuit case law confirms that VMware's construction is correct. *High Tech Med. Instr., Inc. v. New Image Indus., Inc.*, 1997 WL 787052, at *3 (Fed. Cir. Dec. 24, 1997) ("[a] technical term used in a patent document is interpreted as having the meaning that it would be given by persons experienced in the field of the invention, unless it is apparent from the patent and the prosecution history that the inventor used the term with a different meaning.").

And Dr. Madisetti's opinion that "the general concept of a hierarchical token bucket approach . . . was not understood as referring to a specific, particular algorithm or application, but rather the general idea of using e.g., a tiered, tree-like arrangement of token buckets" is contrary to the inventors' own statements about the '818 patent. Madisetti Dec. at ¶ 43. During prosecution the applicants argued the claims were patentable because HTB was different than the general "hierarchical [] structure" and "tree structure" arrangements raised in Dr. Madisetti's

declaration. *See* VMware's Op. Br. at 43. At bottom, IV simply makes a conclusory statement that VMware attempts to "substitute evidence inconsistent with the broad disclosure found in the intrinsic record." But the evidence shows that the inventors intended to claim the specific class-based scheduling algorithm known in the art as the "hierarchical token bucket". *See, e.g.*, Dkt. No. 54 at 52-53 (applicants argued that the claims were patentable because the HTB was different than other "hierarchical or tree structure for storing resource reservations"); *id.* at 51-52 (HTB is referred to as a proper noun or term of art throughout the specification); Dkt. No. 54-2 at 51 (Snoeren Decl. ¶ 111) (the description of HTB in the '818 patent mirrors the description of the well-known HTB algorithm).

B. "enforc[e/ing]", "receiv[e/ing]", "classify[ing]", "compar[e/ing]", "forward[ing]", and "buffer[ing]" ('818 patent claims 1, 17, 30, 32-33, 37-39, 42)

The issue here is whether a patentee can enlarge the scope of a patent beyond what the patentee has clearly defined as "the present invention." VMware's position is that these steps (the "HTB steps") must occur ***in*** the virtual I/O server because that is exactly how the inventors repeatedly and consistently defined the scope of their invention. See, e.g. Dkt. No. 61 at 50 ("The present invention manages QoS of I/O subsystems ***in virtual I/O servers***.") (emphasis added). IV's position is that these steps can be performed anywhere and that the patent is not limited by the inventors' clear statements regarding what they considered to be their invention. As explained in the prior briefing, where, as here, a patentee clearly defines "the present invention" and "repeatedly and consistently" explains the scope of the invention, the claims may not extend any farther. *Id.* at 50-51. Thus, the claims must be limited to methods and systems where these steps occur ***in*** the virtual I/O server as opposed to in the application server or some other part of the system.

IV's response does not change this conclusion. The facts here are nearly identical to the facts in *SkinMedica* and *Microsoft*. *SkinMedica* held that "disclaiming the ordinary meaning of a

claim term—and thus, in effect, redefining it, can be affected through 'repeated and definitive' remarks in the written description." *SkinMedica, Inc. v. Histogen Inc.*, 727 F.3d 1187, 1196 (Fed. Cir. 2013). IV ignores this holding, and like the appellants in *SkinMedica*, argues that the plain meaning is clear despite the distinguishing remarks. *Compare* Dkt. Nos. 60 at 47; 61 at 50 (describing the inventor's remarks about the present invention) with *SkinMedica*, 727 F.3d at 1196. This Court, like *SkinMedica*, should adopt VMware's proposal based on the '818 patent inventor's repeated and definitive remarks about the invention. *See also Microsoft Corp. V. Multi-TechSystems, Inc.*, 357 F.3d 1340, 1346-48 (Fed. Cir. 2004) (construing "sending,' 'transmitting,' and 'receiving' limitations" as "limited to communications over a telephone line and excluding the use of a packet-switched network" based on statements that "broadly describe the overall inventions").

C. "maintaining a connection over a network fabric" ('818 patent claims 1, 17, 30, 32 and 42)

The issue is whether a POSITA would know that a network fabric maintains a connection between physical interfaces. A POSITA would. *See* Ex. 51 (Snoeren Supp. Dec.) at ¶ 2. Indeed, the '818 patent discloses a "virtual I/O server [that] creates virtual device interfaces for application servers 102 to access the I/O subsystems as if the I/O subsystems are directly connected to the application servers." '818 patent at 3:18-21. And the only way for the application server to "access the I/O subsystems" is if the network fabric maintains a connection between two physical interfaces. *Id*.

Rather than address the issue for the Court to decide, IV confuses the issue by citing to, without explanation, the specification's reference to a "virtual [network] interface." IV cites to the disclosure that "virtual network interface, in one implementation, emulates an Ethernet NIC [and] plugs in at the bottom of the network stack." Dkt. No. 60 at 49. But this disclosure is

irrelevant to whether a POSITA would have known that the physical interfaces of the application server and virtual I/O server connect over a network fabric. Ex. 51 (Snoeren Supp. Dec.) at ¶ 6. And IV's next citation is consistent with VMware's position. *See* Dkt. No. 60 at 48 ("[i]nserted into the network and storage protocol stacks are virtual interface drivers configured to intercept storage and network I/O messages") (citing '818 patent at 4:2-4); *see also* Ex. 51 (Snoeren Supp. Dec.) at ¶ 5. IV cites no intrinsic support of network fabric connecting interfaces other than physical interfaces.

IV's remaining arguments are misplaced. IV mischaracterizes VMware is not reading in embodiment limitations. Rather, VMware's proposal is that "the written description's only use of the [disputed term] supports [VMware's] construction." *Apple Inc. v. Andrea Elec. Corp.*, 949 F.3d 697, 707 (Fed. Cir. Feb. 2020); *see also* Ex. 51 (Snoeren Supp. Dec.) at ¶ 2. Nor does Dr. Snoeren's discussion of the only implementation—that the network fabric connects physical interfaces—described in the patent contradict VMware's position. And IV ignores the other examples of this. *See* Dkt. No. 54-2 at 52 (Snoeren Decl. at ¶114)) and the numerous other examples showing the only implementation disclosed in the '818 patent. Dkt. No. 61 at 52-57. Finally, IV tries to construe the disputed term without the phrase "over a network fabric." Dkt. No. 60 at 48. But IV fails to acknowledge that 1) the endpoints of the network fabric is where the connection is maintained, and 2) "to a virtual [storage/network] interface layer of an application server" is merely referring to Figure 2, which shows how the connection is made over the network fabric to the "virtual [storage/network] interface layer." Ex. 51. (Snoeren Supp. Dec.) at ¶¶ 3-4.

At bottom, VMware's proposal is the only way a POSITA could read the term and IV's proposal attempts to broaden the scope under the guise of a "plain and ordinary meaning" construction. This Court should grant VMware's proposed construction.

D. "virtual storage network interface layer of an application server" / "virtual network interface layer of an application server"/ "virtual interface layer of an application server" ('818 patent claims 1, 17, 30, 32 and 42)[20]

IV's response focuses on the construction VMware proposed prior to knowing the plain and ordinary meaning of IV's proposal. Specifically, IV focuses on the terms ". . . to higher layers" and "virtual node." Dkt. No. 60 at 50-51. VMware's proposal, based on the meaning IV proposed for the first time in its opening brief, does not include the terms ". . . to higher layers" and "virtual node." Instead, VMware's proposal is consistent with IV's explanation of the plain and ordinary meaning of the disputed term provided in its opening brief. Dkt. No. 61 at 54-55. For the reasons stated in VMware's response brief, VMware's proposal should be adopted.

E. "one or more input/output virtualization modules comprising computer-readable instructions operative to cause the one or more processors to" performs functions terms ('818 patent claim 17)

The issue here is whether "input/output virtualization module" is a "nonce term" that invokes means plus function law, and if so, whether the patent discloses sufficient structure. IV reliance on its expert declaration is misplaced—that declaration simply offers conclusory statements without providing substantive explanation. Dkt. No. 60 at 51. The term "input/output virtualization module" is not a reference to a conventional structure which is why IV failed to provide any substantive evidence to the contrary.[21] IV's remaining arguments as to the alleged corresponding structure for the disputed terms are misplaced. *See, e.g.,* Dkt. No. 61 at 56-57.

VI. CONCLUSION

For the reasons stated herein, VMware respectfully requests the Court adopt its proposed constructions for the disputed terms and phrases.

---

[20] As mentioned in VMware's response it's proposal is "interface layers (e.g., virtual network interface 220, virtual HBA 208a) that emulate layers of a networking or storage protocol stack."
[21] IV's citation to *Zeroclick, LLC* is inapplicable here. *See, e.g.,* Dkt. No. 61 at 56.

Dated: April 10, 2020

Respectfully submitted,

WINSTON & STRAWN LLP

/s/ *Katherine Vidal*

Katherine Vidal
Admitted *Pro Hac Vice*
KVidal@winston.com
Michael R. Rueckheim
Texas State Bar No. 24081129
MRuekheim@winston.com
WINSTON & STRAWN LLP
275 Middlefield Road, Suite 205
Menlo Park, CA 94025
Telephone: (650) 858-6500
Facsimile: (650) 858-6559

Thomas M. Melsheimer
Texas State Bar No. 13922550
M. Brett Johnson
Texas State Bar No. 00790975
MBJohnson@winston.com
Michael A. Bittner
MBittner@winston.com
Admitted *Pro Hac Vice*
Winston & Strawn LLP
2121 N. Pearl St., 9th Floor
Dallas, TX 75201
Telephone: (214) 453-6500
Facsimile: (214) 453-6400

Vivek V. Krishnan
Admitted *Pro Hac Vice*
VKrishnan@winston.com
DaWanna McCray
Admitted *Pro Hac Vice*
DMcCray@winston.com
Winston & Strawn LLP
35 W. Wacker Drive
Chicago, IL 60601
Telephone: (312) 558-5600
Facsimile: (312) 558-5700

William M. Logan
Texas State Bar No. 24106214
WLogan@winston.com
Winston & Strawn LLP
1111 Louisiana Street
Houston, TX 77002
Telephone: (713) 651-2766
Facsimile: (713) 651-2700

ATTORNEYS FOR DEFENDANT VMWARE, INC.

CERTIFICATE OF SERVICE

I hereby certify that, on April 10, 2020, the foregoing document was electronically filed with the Clerk of Court using the Court's CM/ECF system which will send notification of such filing to all counsel of record, including counsel of record for Plaintiffs Intellectual Ventures I LLC and Intellectual Ventures II LLC.

/s/ *Katherine Vidal*
Katherine Vidal